



US008949414B2

(12) **United States Patent**
Raja et al.

(10) **Patent No.:** **US 8,949,414 B2**
(45) **Date of Patent:** **Feb. 3, 2015**

(54) **SYSTEMS AND METHODS FOR SCALABLE
N-CORE STATS AGGREGATION**

(75) Inventors: **Murali Raja**, Santa Clara, CA (US);
Sridhar Srinivasan, Bangalore, CA
(US); **Saravana Annamalaisami**,
Bangalore, CA (US); **Devesh Prakash**,
Bangalore, CA (US); **Josephine**
Suganthi, Sunnyvale, CA (US); **Anil**
Shetty, Union City, CA (US)

(73) Assignee: **Citrix Systems, Inc.**, Fort Lauderdale,
FL (US)

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 375 days.

(21) Appl. No.: **13/337,706**

(22) Filed: **Dec. 27, 2011**

(65) **Prior Publication Data**

US 2012/0226804 A1 Sep. 6, 2012

Related U.S. Application Data

(60) Provisional application No. 61/428,124, filed on Dec.
29, 2010.

(51) **Int. Cl.**
G06F 15/173 (2006.01)
H04L 12/26 (2006.01)
H04L 29/06 (2006.01)

(52) **U.S. Cl.**
CPC **H04L 43/028** (2013.01); **H04L 29/06**
(2013.01); **H04L 12/26** (2013.01)
USPC **709/224**; **709/223**

(58) **Field of Classification Search**
USPC **709/224**
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,539,501 B1	3/2003	Edwards	
7,447,209 B2 *	11/2008	Jeffay et al.	370/392
7,467,202 B2 *	12/2008	Savchuk	709/224
7,526,807 B2 *	4/2009	Chao et al.	726/23
7,610,330 B1 *	10/2009	Quinn et al.	709/201
2002/0001307 A1 *	1/2002	Nguyen et al.	370/386
2002/0016937 A1 *	2/2002	Houh	714/43
2002/0035628 A1 *	3/2002	Gil et al.	709/224
2002/0065938 A1 *	5/2002	Jungck et al.	709/246

(Continued)

OTHER PUBLICATIONS

Deri and Fusco, "Exploiting commodity multicore systems for net-
work traffic analysis" IMC'10, ACM, Nov. 2010, pp. 218-224.*

(Continued)

Primary Examiner — David Lazaro

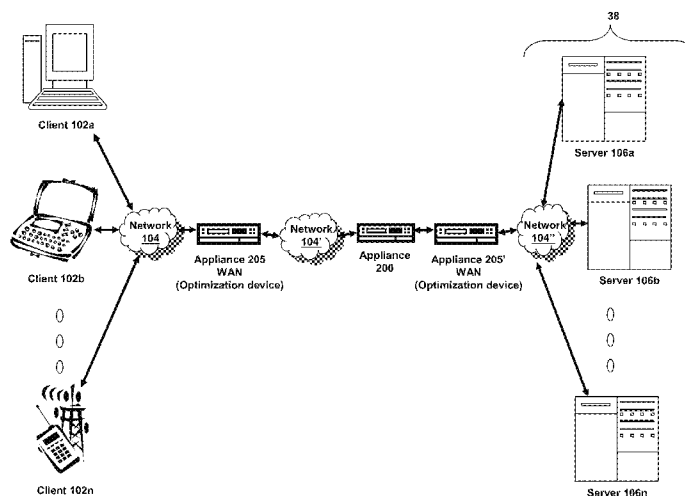
Assistant Examiner — Robert Shaw

(74) *Attorney, Agent, or Firm* — Foley & Lardner LLP;
Christopher J. McKenna

(57) **ABSTRACT**

The present invention is directed towards systems and meth-
ods for aggregating and providing statistics from cores of a
multi-core system intermediary between one or more clients
and servers. The system may maintain in shared memory a
global device number for each core of the multi-core system.
The system may provide a thread for each core of the multi-
core system to gather data from the corresponding core. A
first thread may generate aggregated statistics from a corre-
sponding core by parsing the gathered data from the corre-
sponding core. The first thread may transfer the generated
statistics to a statistics log according to a schedule. The sys-
tem may adaptively reschedule the transfer by monitoring the
operation of each computing thread. Responsive to a request
from a client, an agent of the client may obtain statistics from
the statistics log.

20 Claims, 21 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

2002/0143929	A1 *	10/2002	Maltz et al.	709/224	2010/0180101	A1 *	7/2010	Trumler et al.	712/30
2003/0056200	A1 *	3/2003	Li et al.	717/128	2010/0251160	A1 *	9/2010	Shafi et al.	715/772
2003/0145077	A1 *	7/2003	Khan et al.	709/224	2010/0322089	A1	12/2010	Raja et al.	
2004/0210320	A1 *	10/2004	Pandya	700/1	2011/0161618	A1 *	6/2011	Bellows et al.	711/202
2005/0108518	A1 *	5/2005	Pandya	713/151	2011/0206064	A1 *	8/2011	Head et al.	370/474
2005/0111367	A1 *	5/2005	Jonathan Chao et al.	370/235	2011/0296024	A1 *	12/2011	Madani et al.	709/226
2005/0276275	A1 *	12/2005	Sabry et al.	370/419	2012/0022832	A1 *	1/2012	Shannon et al.	702/189
2006/0004942	A1 *	1/2006	Hetherington et al.	711/3	2012/0072575	A1 *	3/2012	King et al.	709/224
2006/0020715	A1 *	1/2006	Jungck	709/246					
2006/0083231	A1 *	4/2006	Jeffay et al.	370/389					
2006/0143333	A1 *	6/2006	Minturn et al.	710/52					
2006/0179156	A1 *	8/2006	Eatherton et al.	709/238					
2006/0212677	A1 *	9/2006	Fossum	712/1					
2006/0259831	A1 *	11/2006	Sohm et al.	714/45					
2007/0006293	A1 *	1/2007	Balakrishnan et al.	726/13					
2007/0081471	A1 *	4/2007	Talley et al.	370/252					
2007/0083644	A1 *	4/2007	Miller et al.	709/224					
2007/0248029	A1 *	10/2007	Merkey et al.	370/255					
2007/0282572	A1 *	12/2007	Larus	703/2					
2007/0294395	A1 *	12/2007	Strub et al.	709/224					
2007/0294671	A1 *	12/2007	Demetriou et al.	717/124					
2008/0126750	A1 *	5/2008	Sistla	712/30					
2008/0250213	A1 *	10/2008	Holt	711/159					
2008/0253395	A1 *	10/2008	Pandya	370/469					
2008/0262765	A1 *	10/2008	Sedeh	702/67					
2009/0019538	A1 *	1/2009	Pandya	726/13					
2009/0092057	A1 *	4/2009	Doctor et al.	370/252					
2009/0164993	A1 *	6/2009	Flake et al.	718/1					
2009/0187713	A1 *	7/2009	Zedlewski et al.	711/130					
2009/0313445	A1 *	12/2009	Pandey et al.	711/162					

OTHER PUBLICATIONS

Dashtbozorgi and Azgomi a high-performance and scalable multi-core aware software solution for network monitoring, Springer, Aug. 17, 2010.*

Deri and Fusco, "Exploiting commodity multicore systems for network traffic analysis" IMC'10, ACM, Nov. 2010, pp. 218-224 <<http://web.archive.org/web/20100301113400/http://luca.ntop.org/MulticorePacketCapture.pdf>>.*

H3c.com, Multi-Core Distributed NetStream Technology White Paper, Nov. 29, 2008.*

Kahn et al, A Programmable Architecture for Scalable and Real-time Network Traffic Measurements, ANCS'08 ACM (2008) pp. 109-118.*

Wun et al, Parallelization of Snort on a Multi-core Platform, ANCS'09, ACM (2009), pp. 173-174.*

Roesch, Snort: Lightweight intrusion detection for networks.

USENIX LISA Systems Administration Conference, Nov. 1999.*

International Preliminary Report on Patentability for PCT/US2011/067380 dated Jul. 11, 2013.

International Search Report issued Mar. 16, 2012 in PCT Application No. PCT/US2011/067380.

* cited by examiner

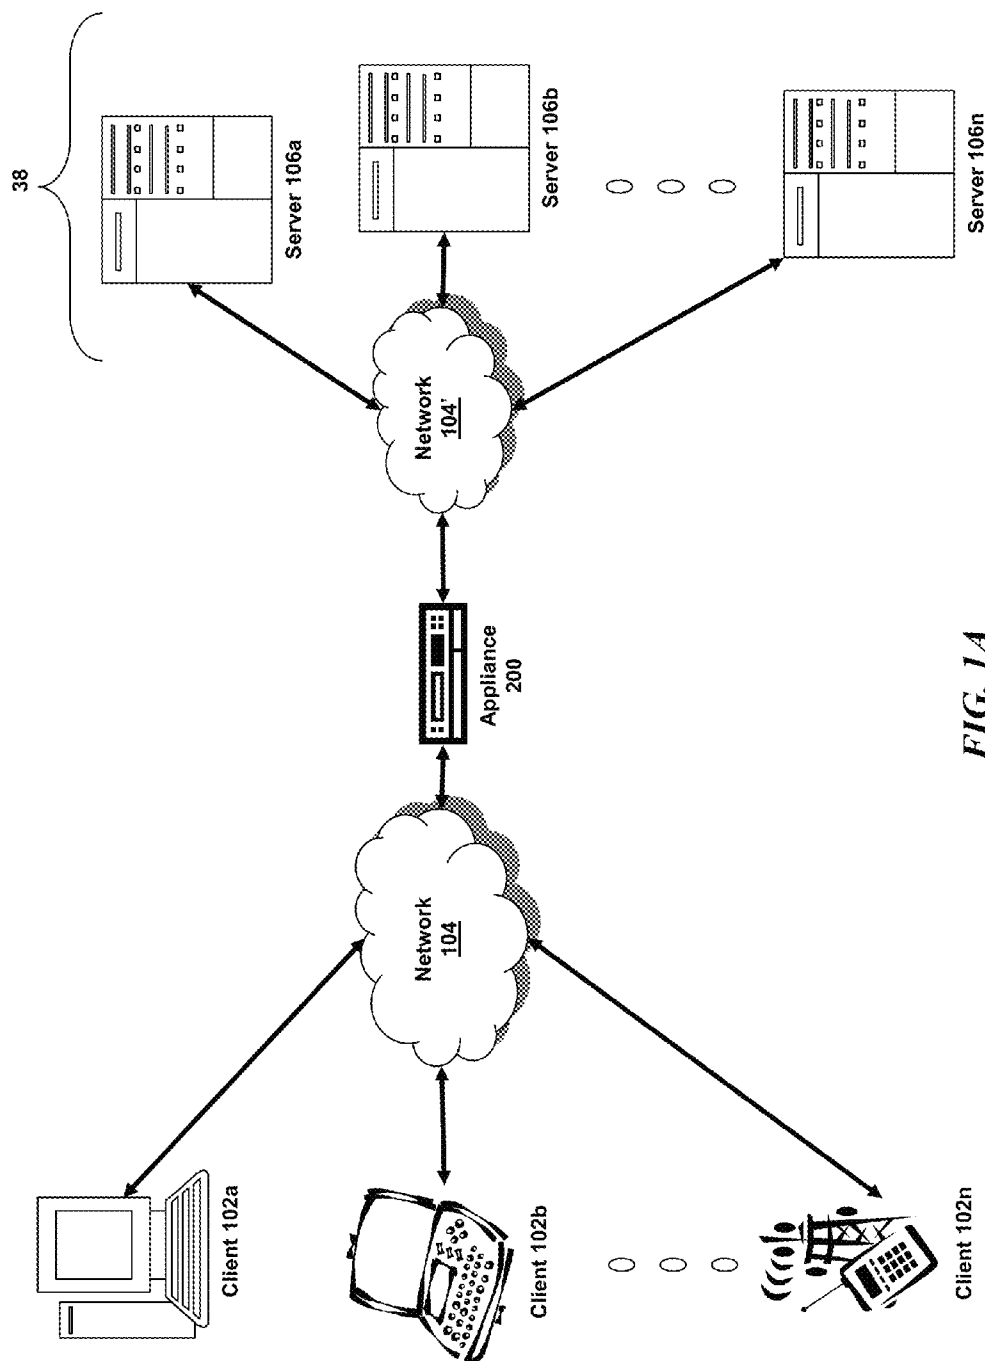
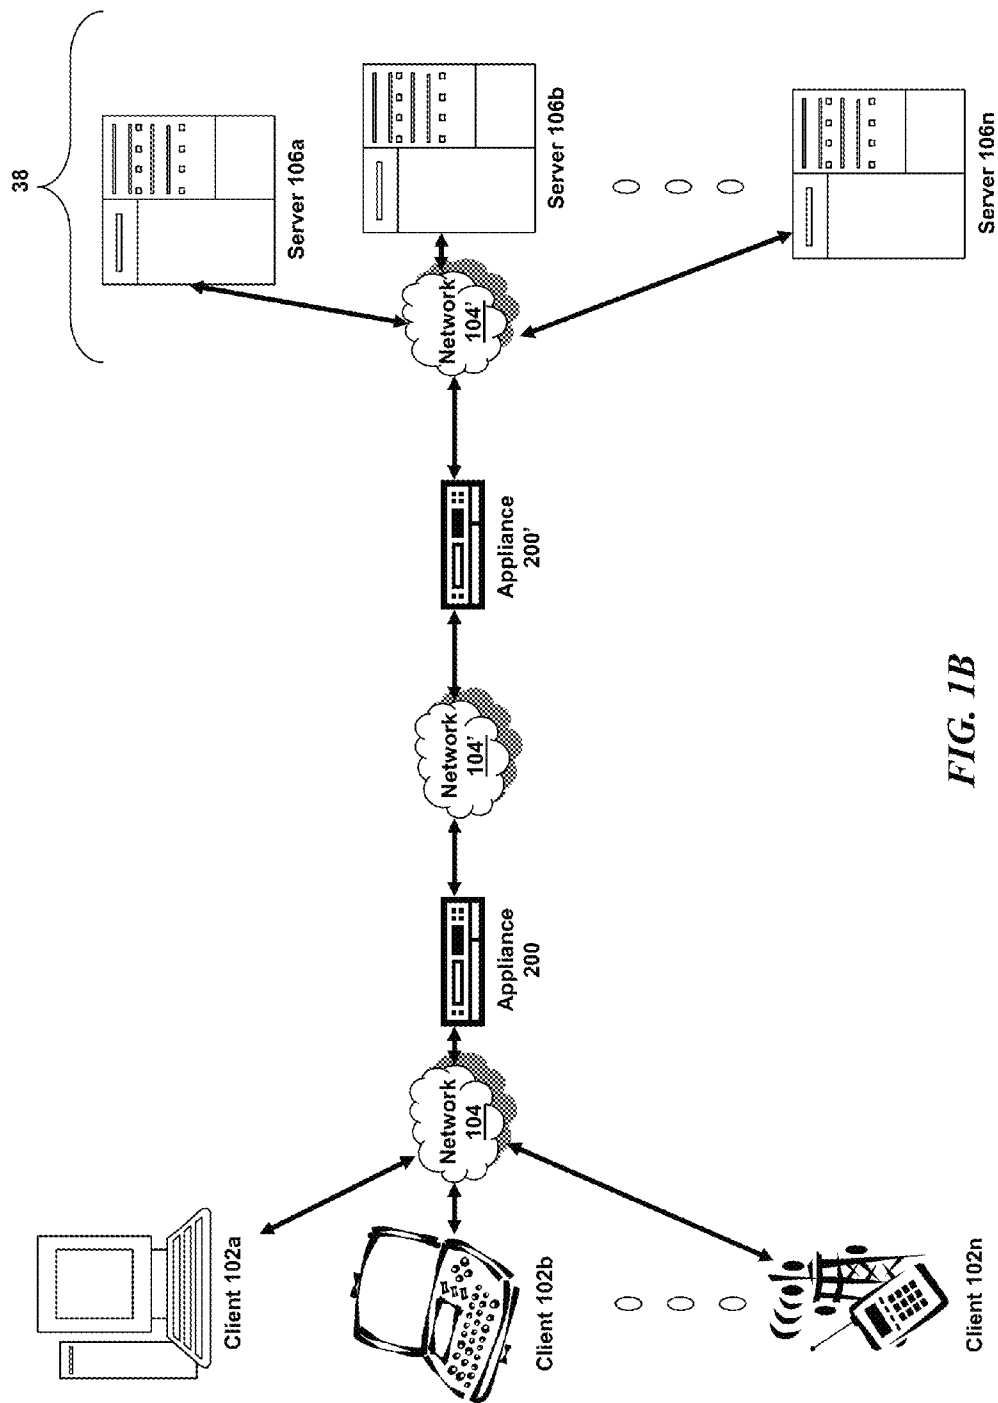


FIG. 1A



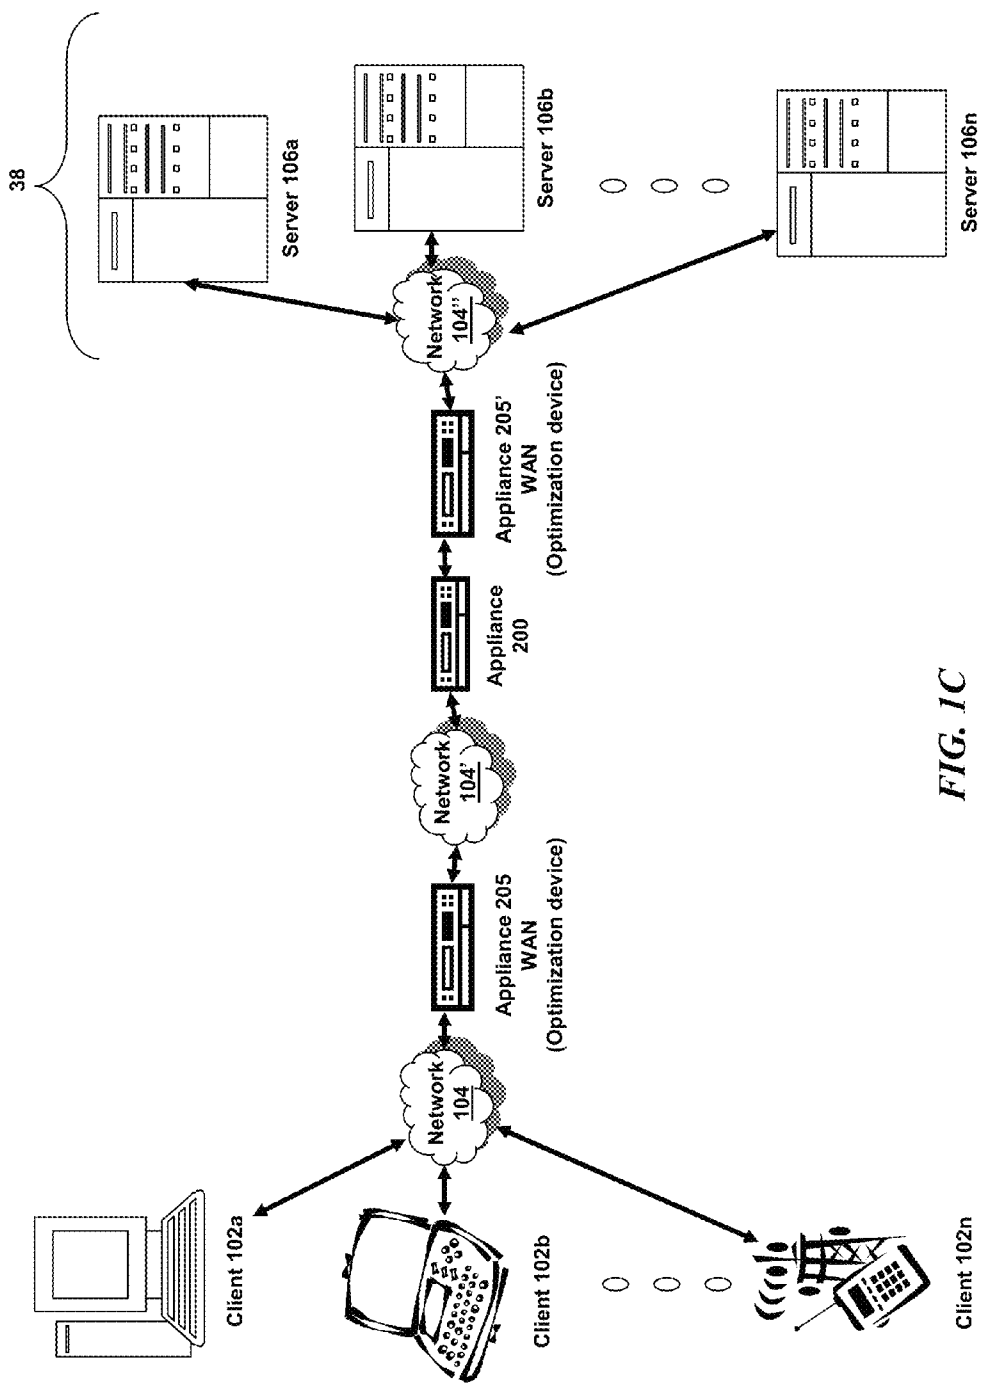


FIG. 1C

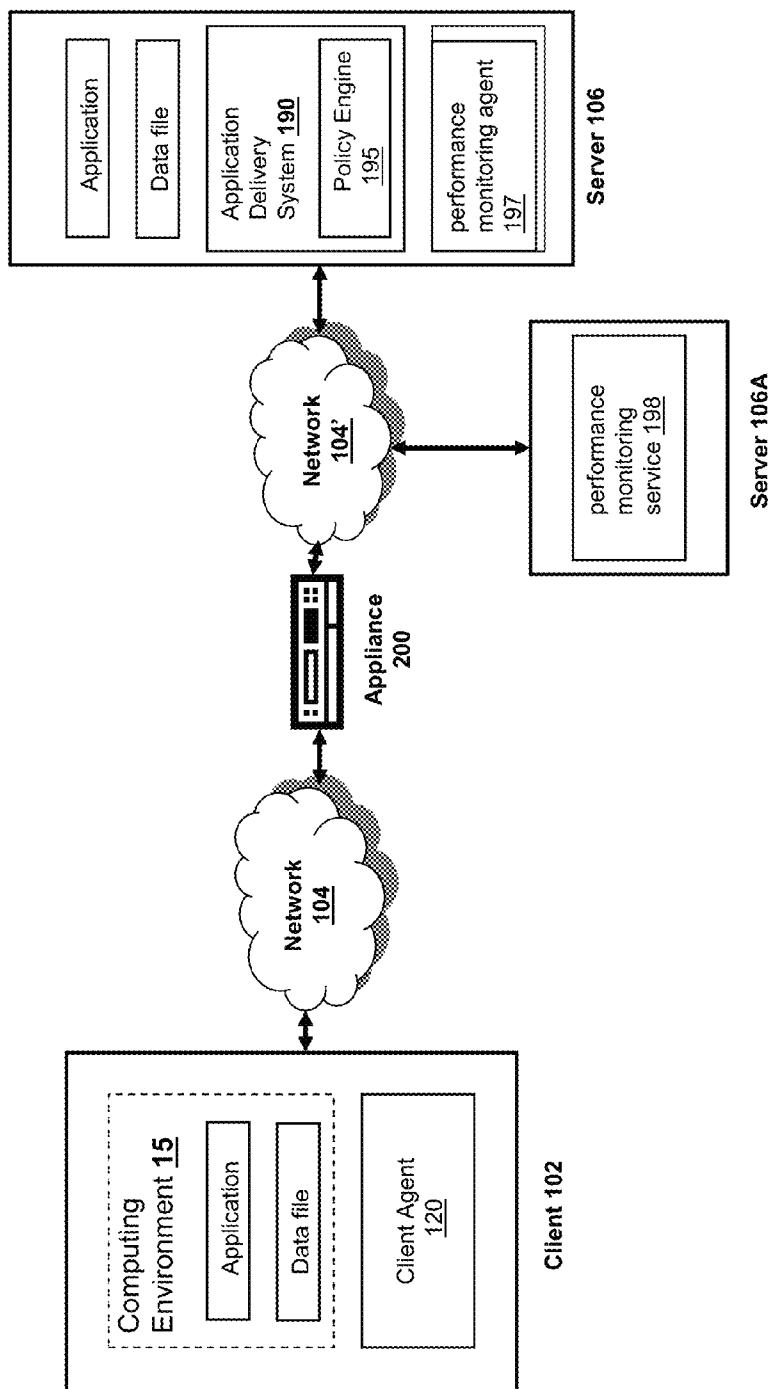


FIG. 1D

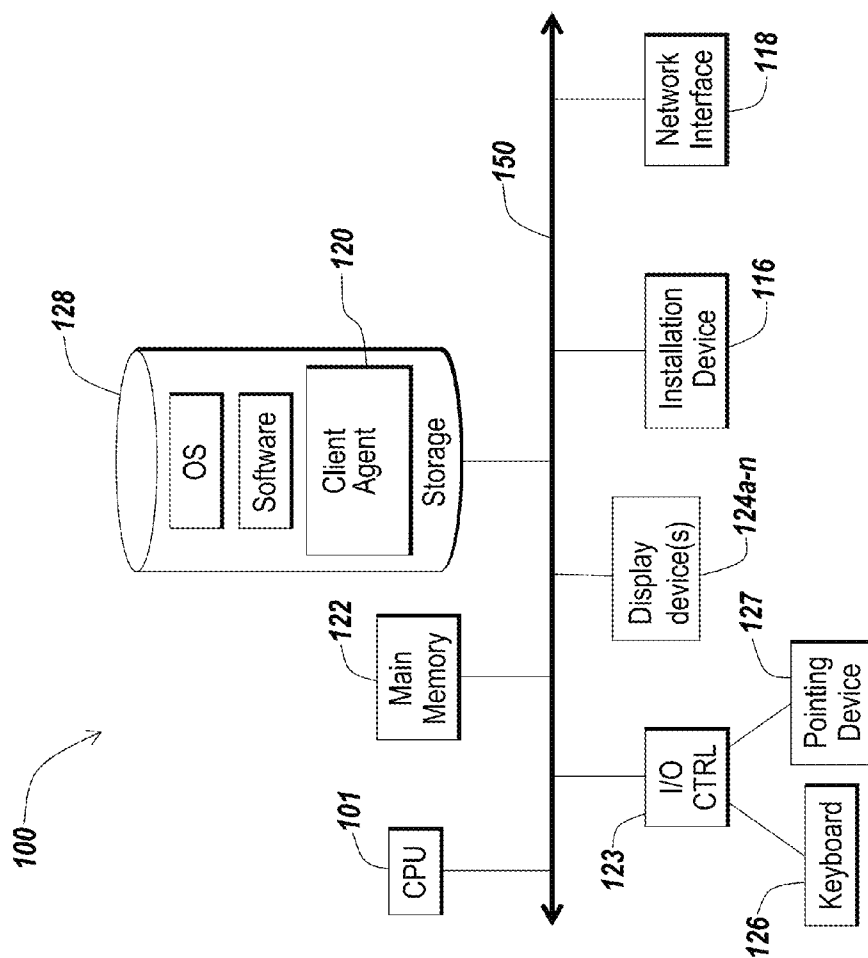
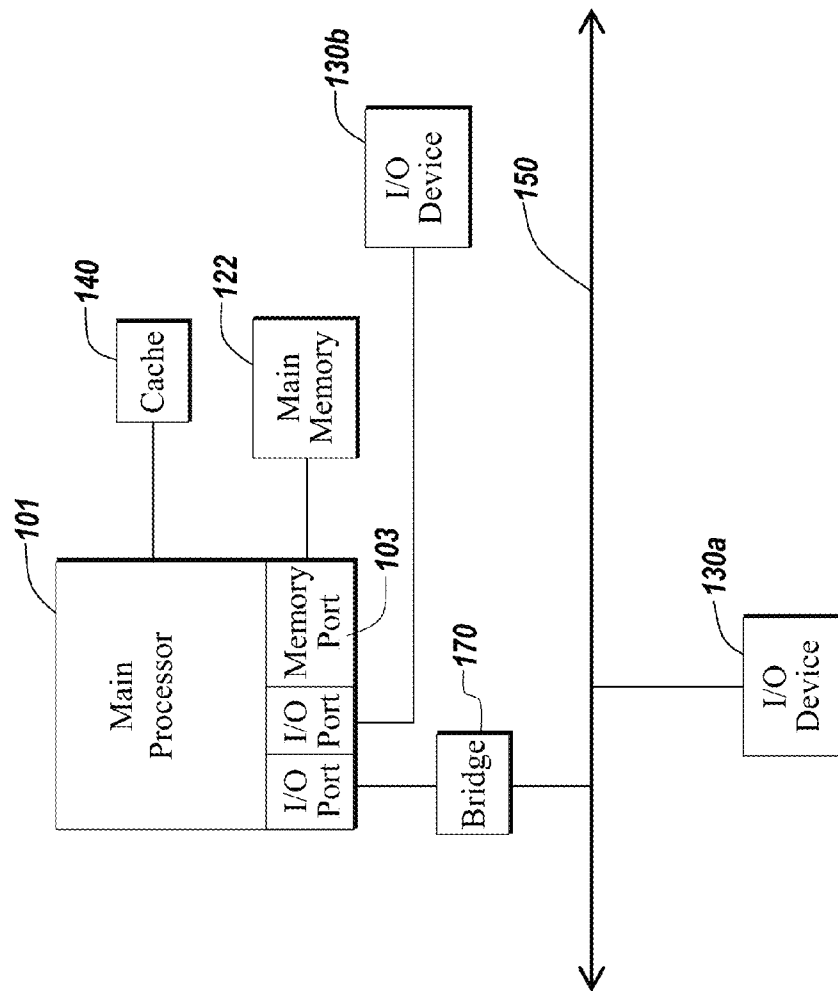


FIG. 1E

**FIG. 1F**

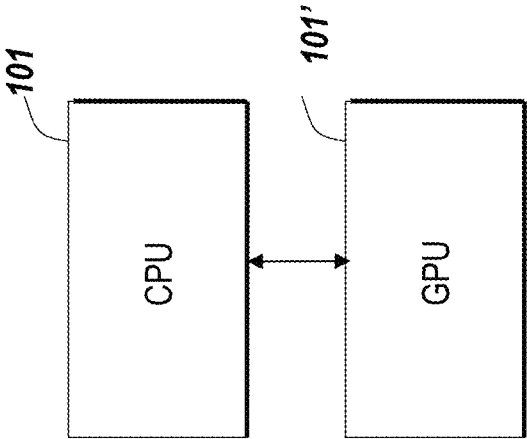


FIG. 1H

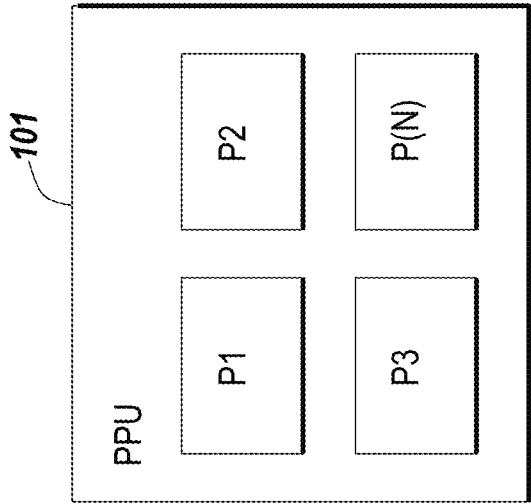


FIG. 1G

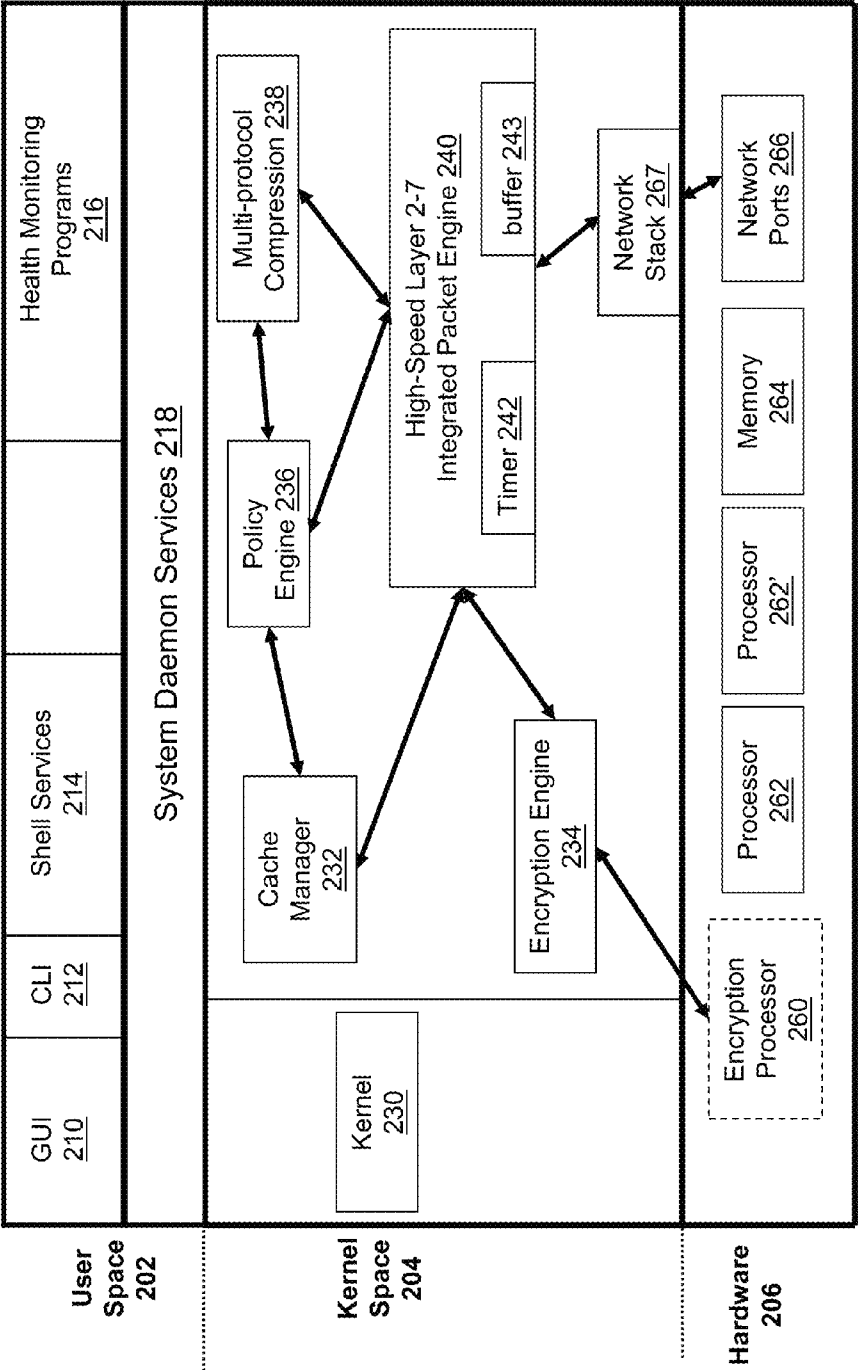


FIG. 2A

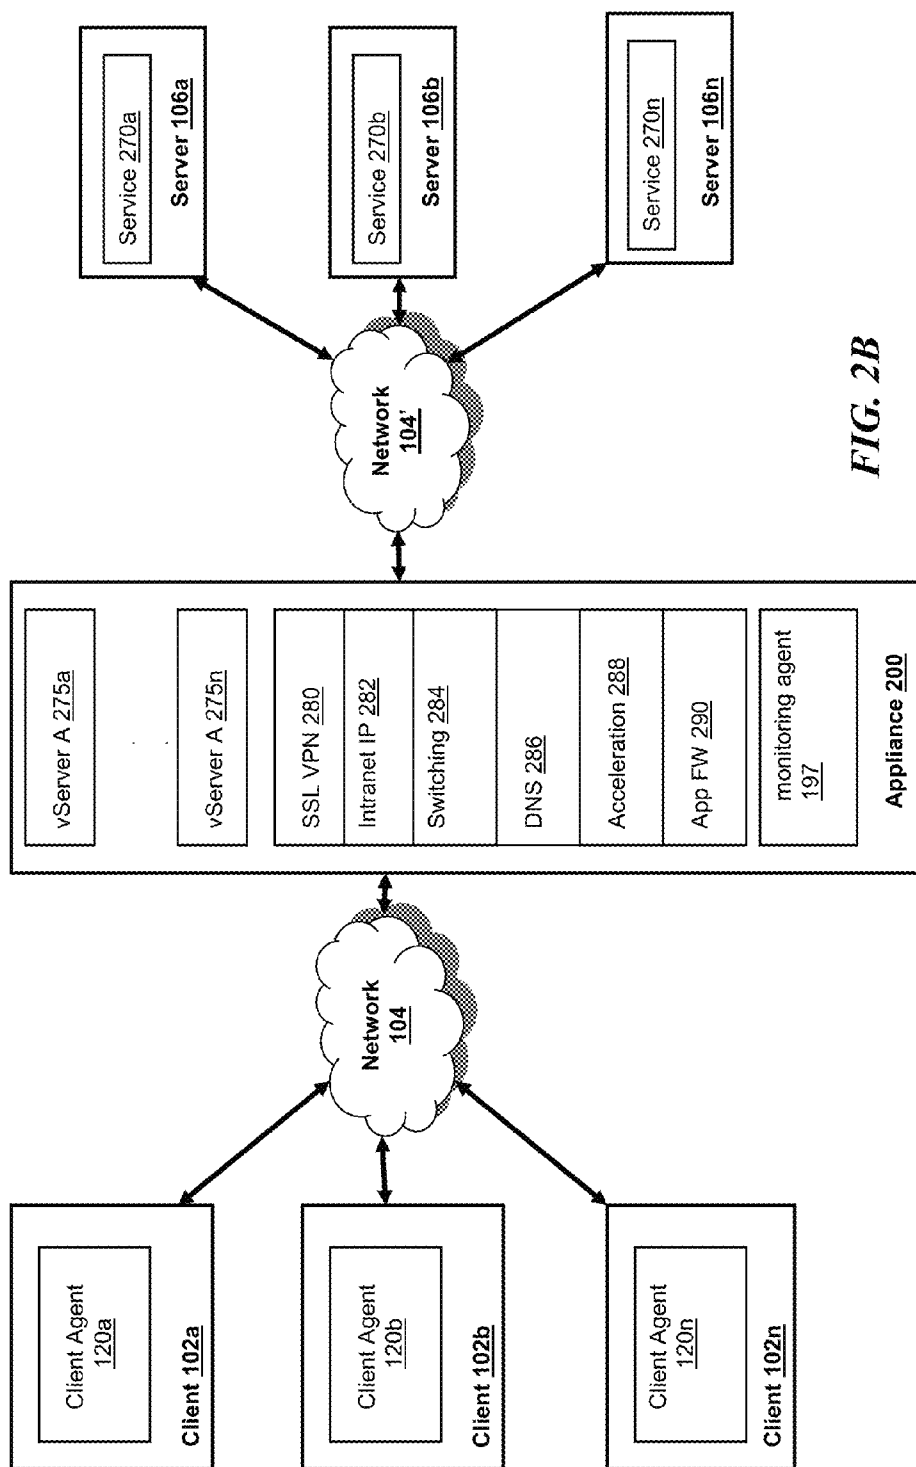


FIG. 2B

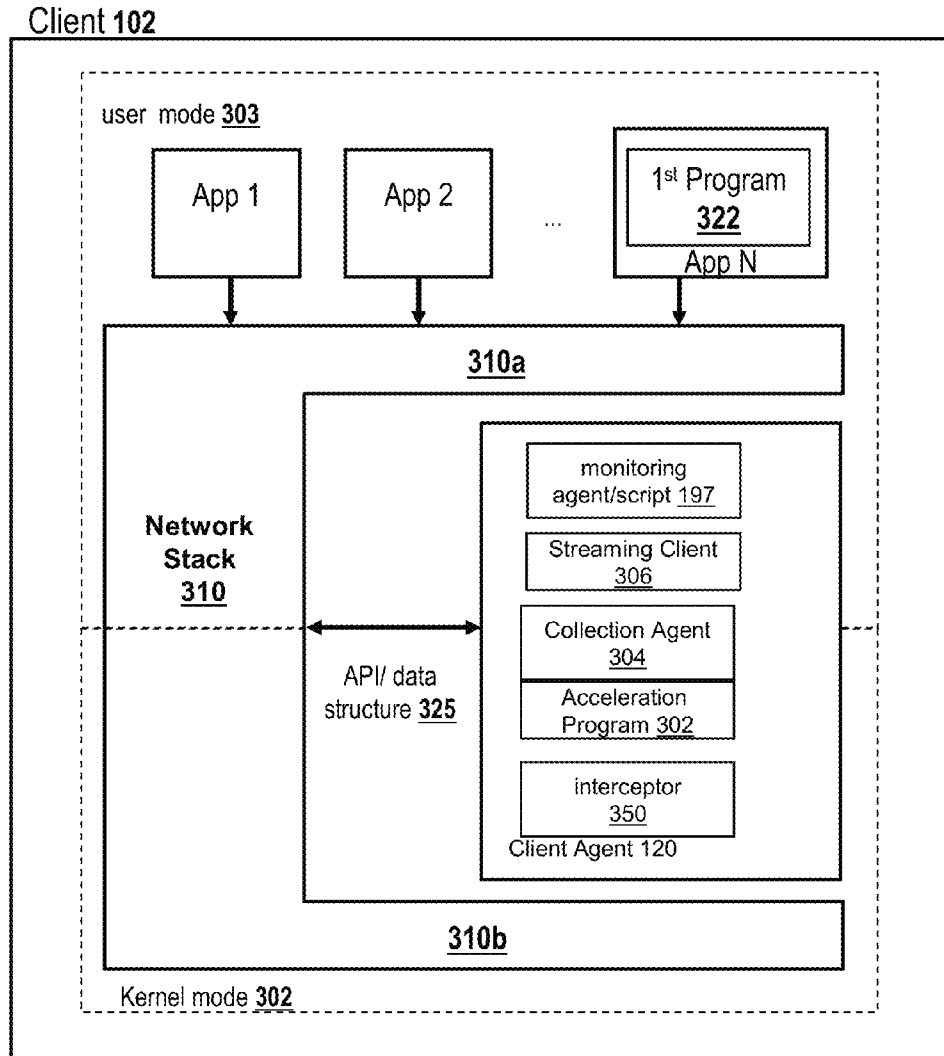


FIG. 3

100

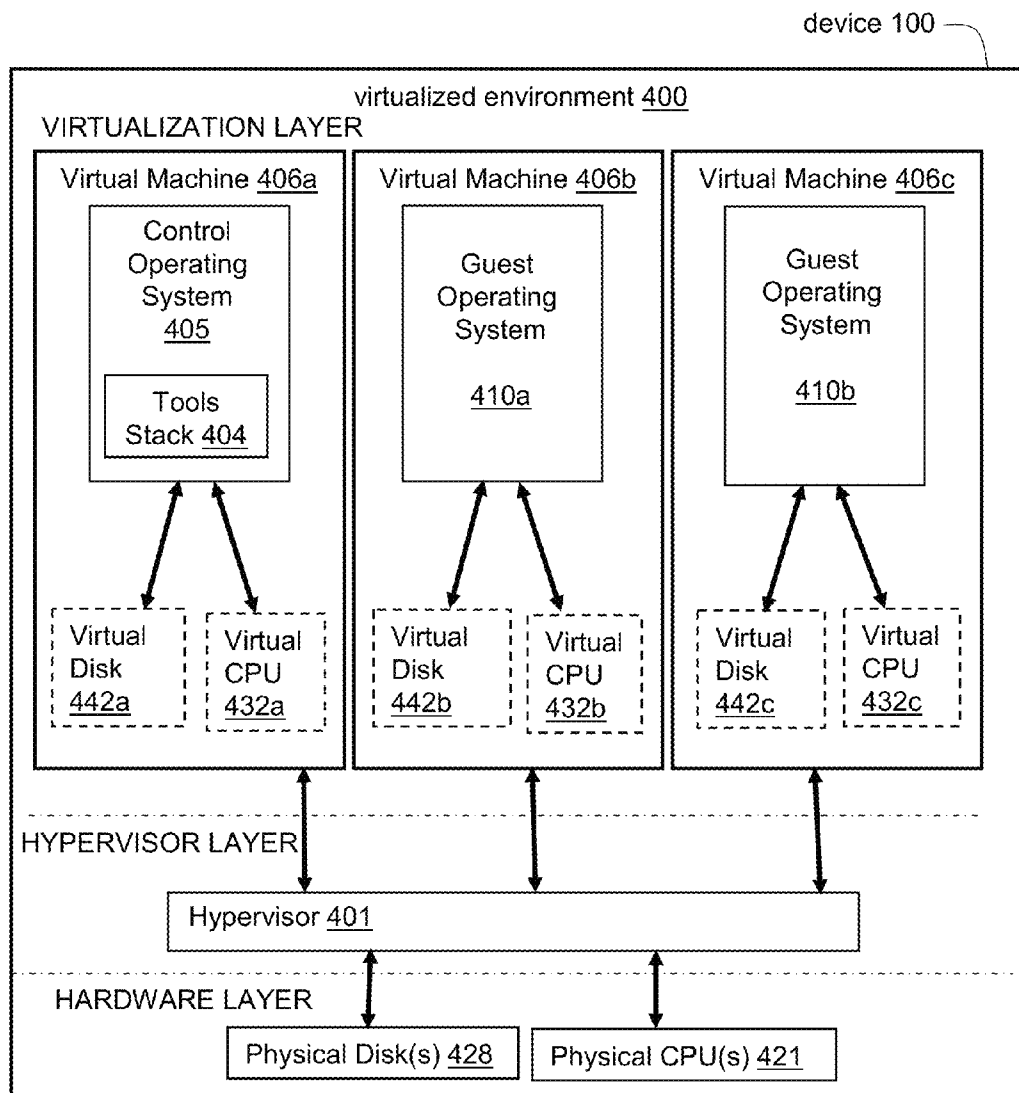
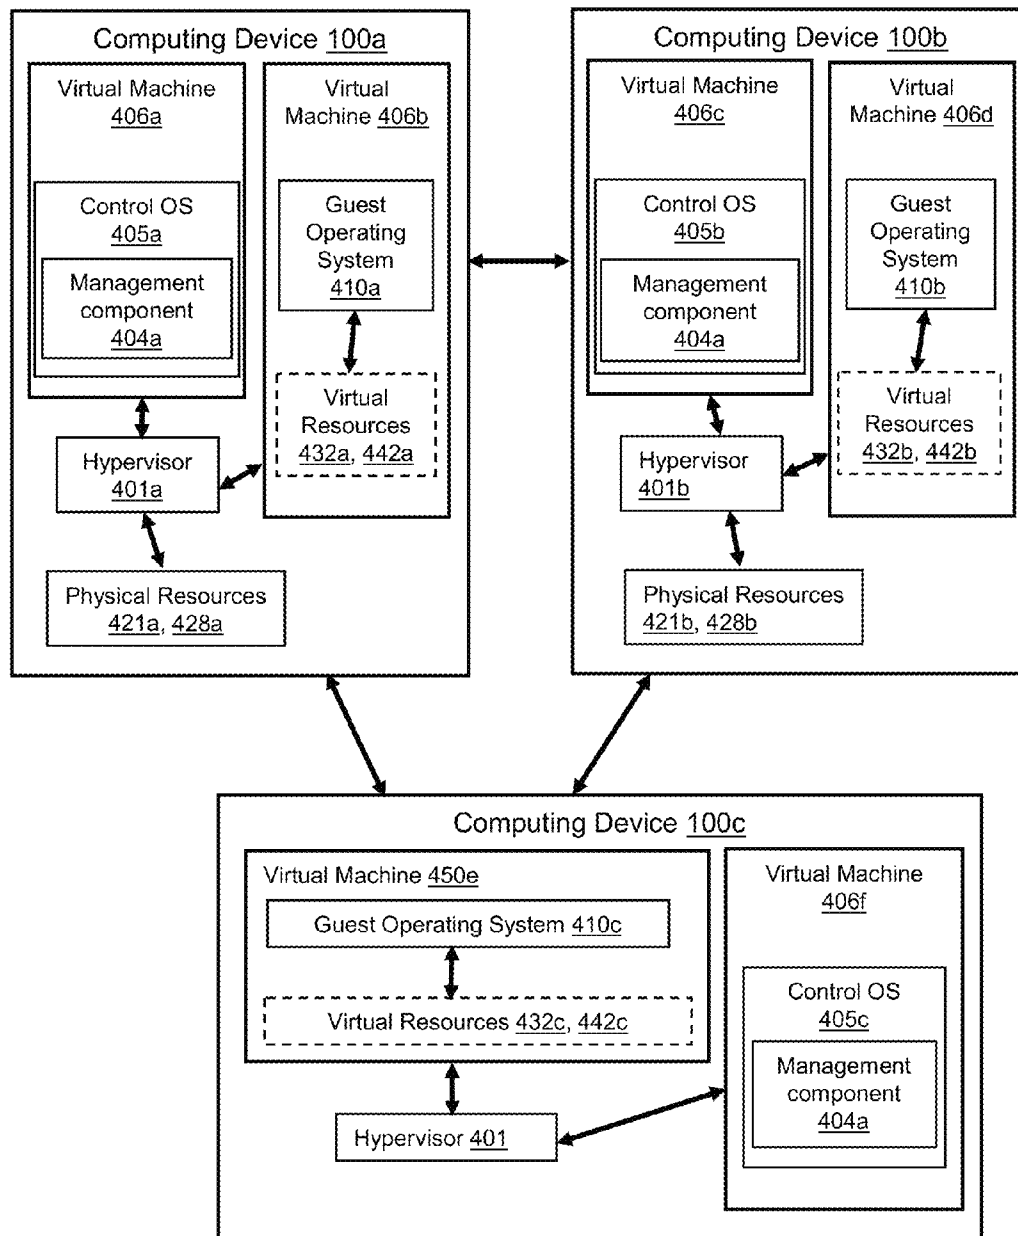
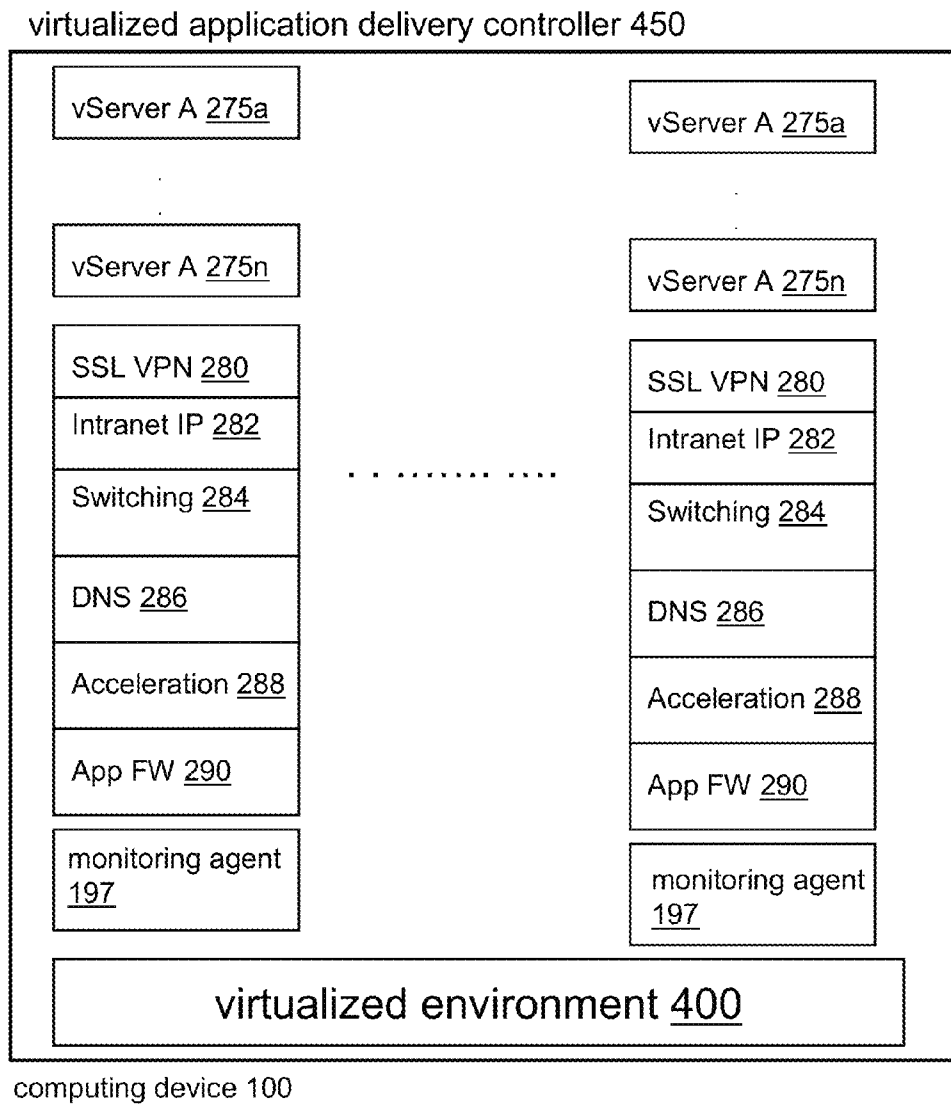


FIG. 4A

**FIG. 4B**

**FIG. 4C**

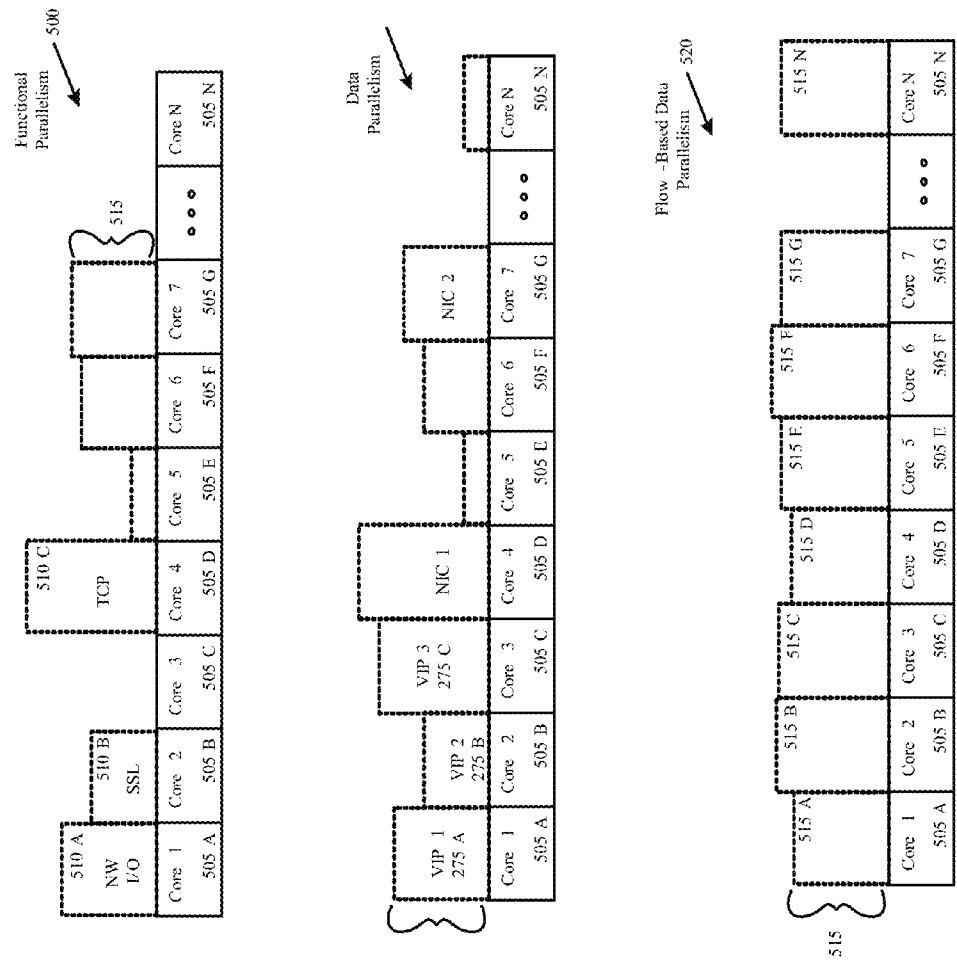


FIG. 5A

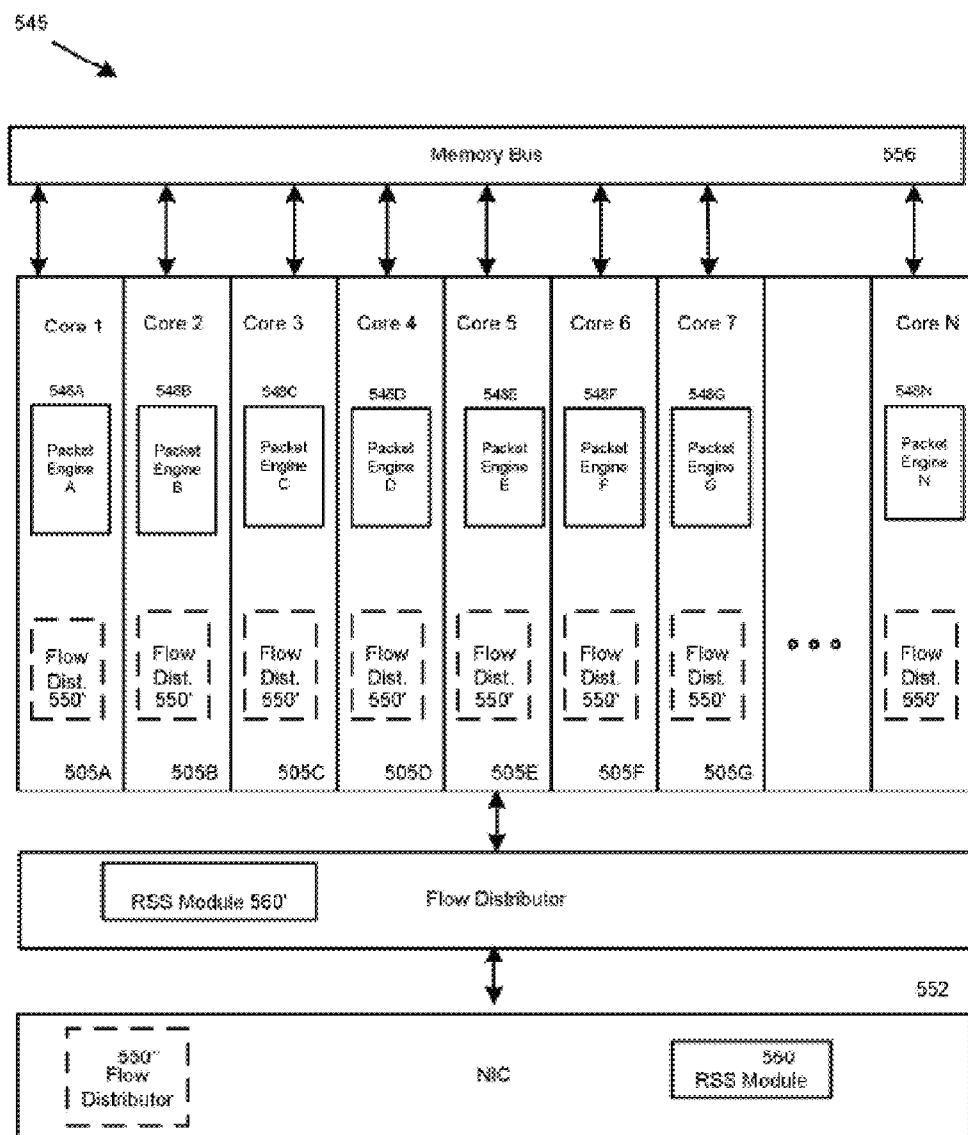


FIG. 5B

575
↗

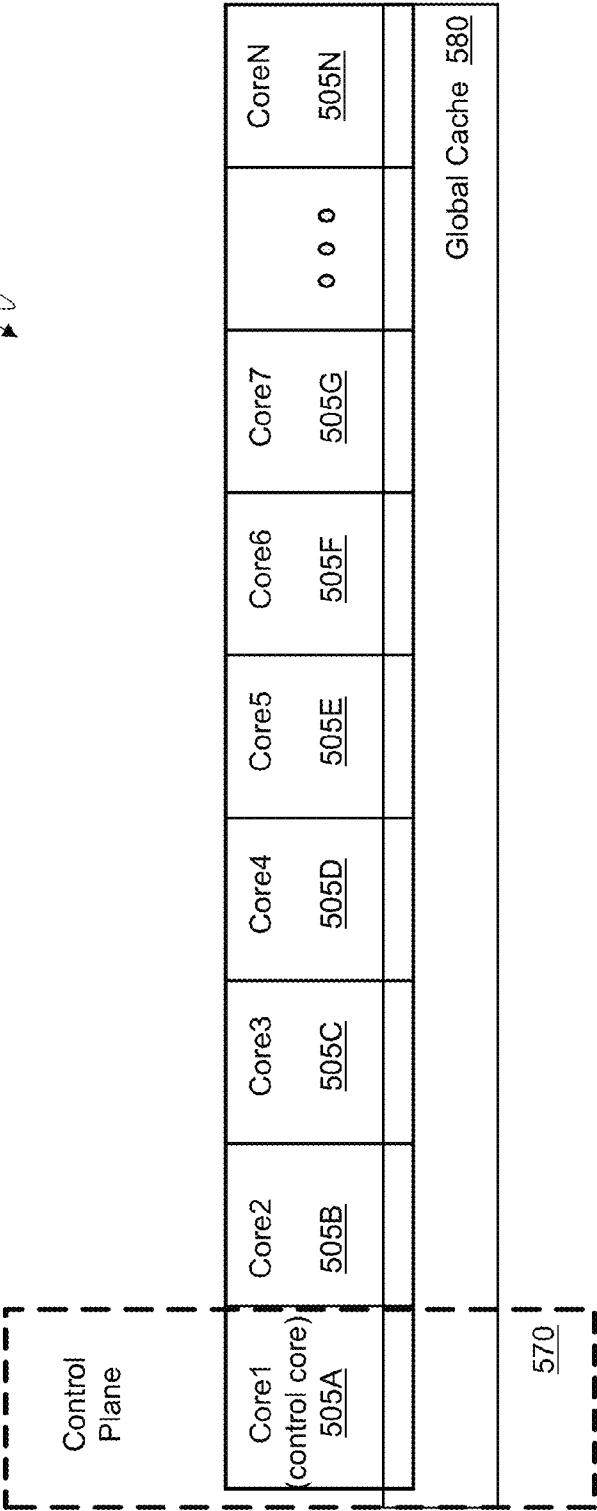


FIG. 5C

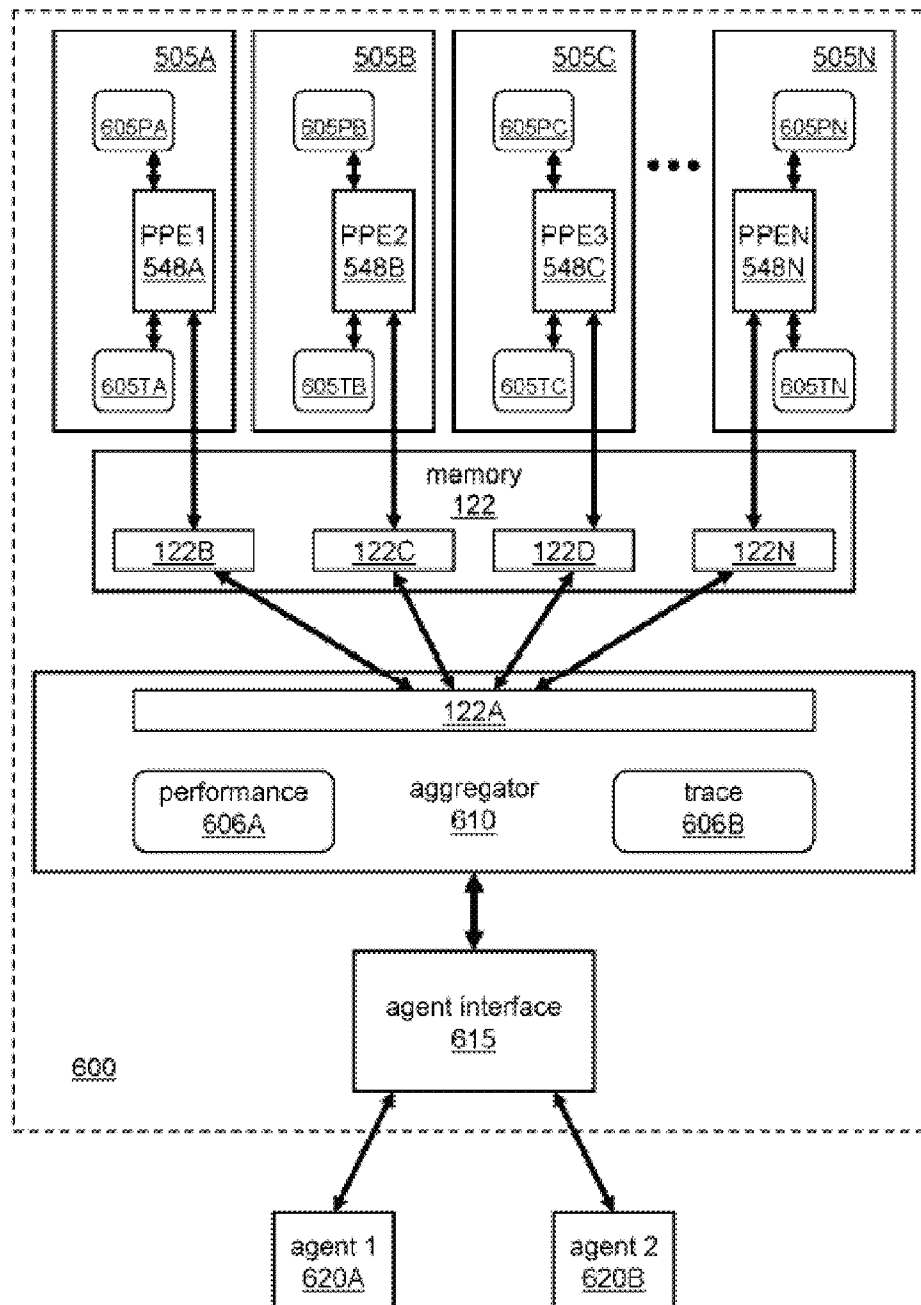
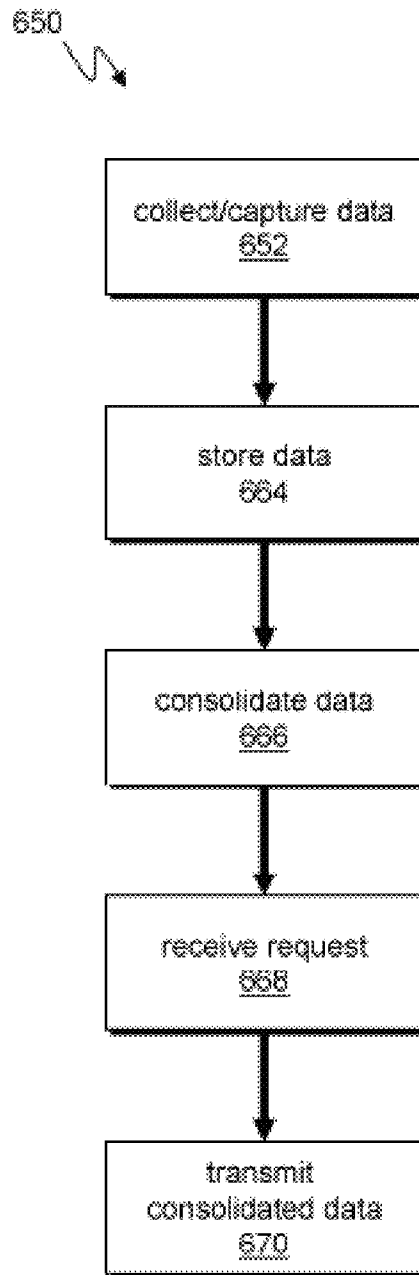


FIG. 6A

**FIG. 6B**

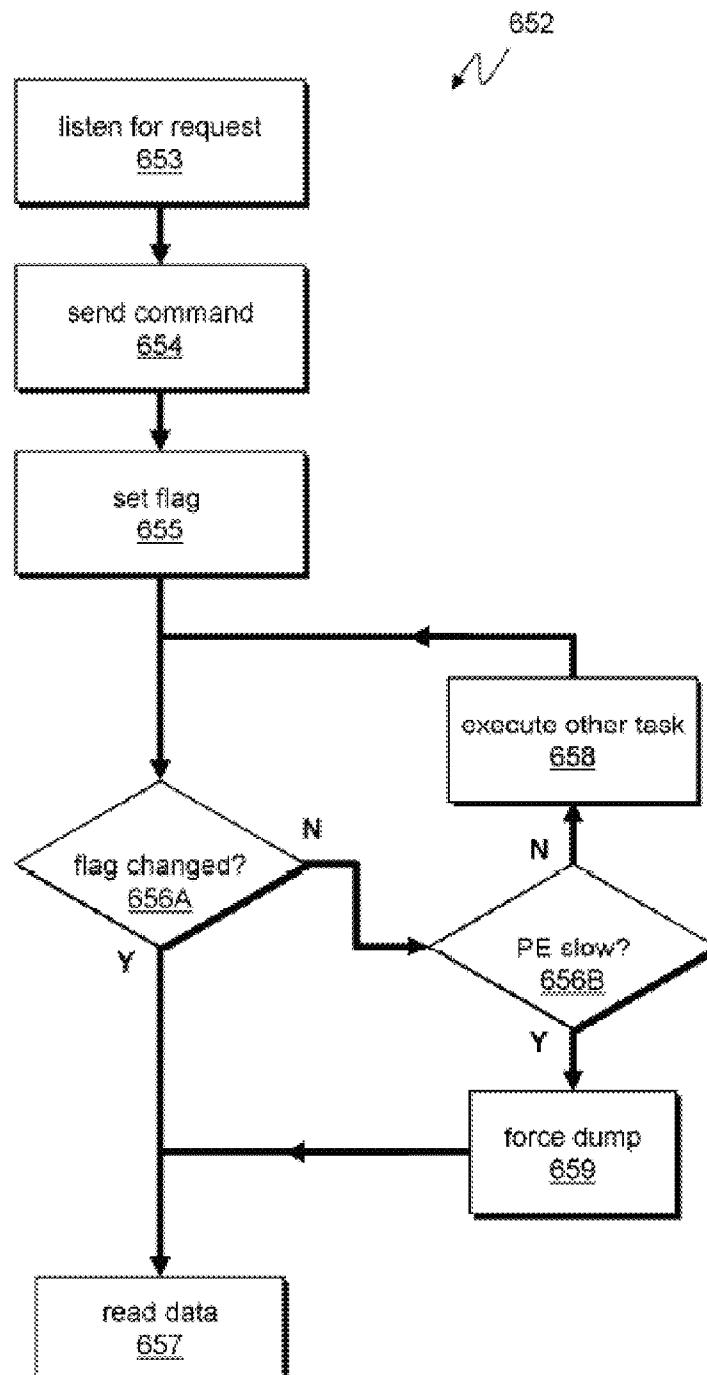


FIG. 6C

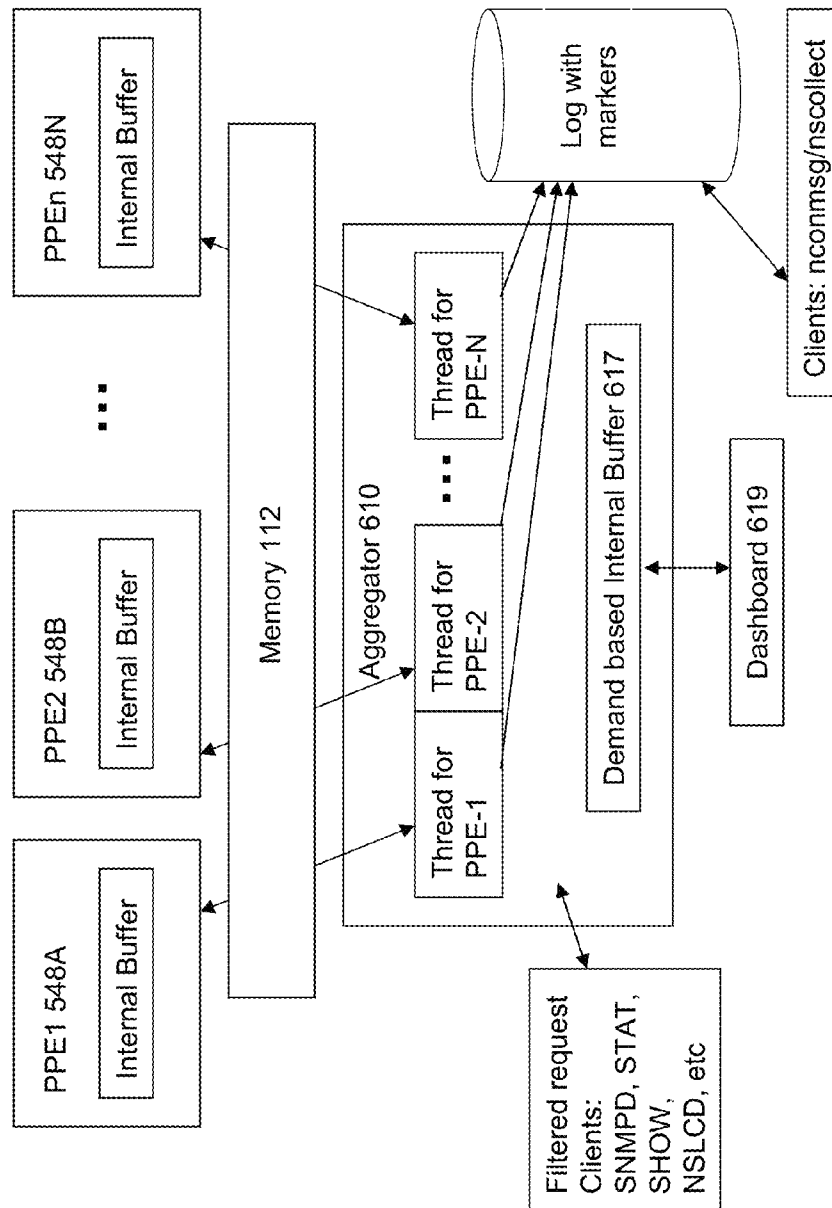
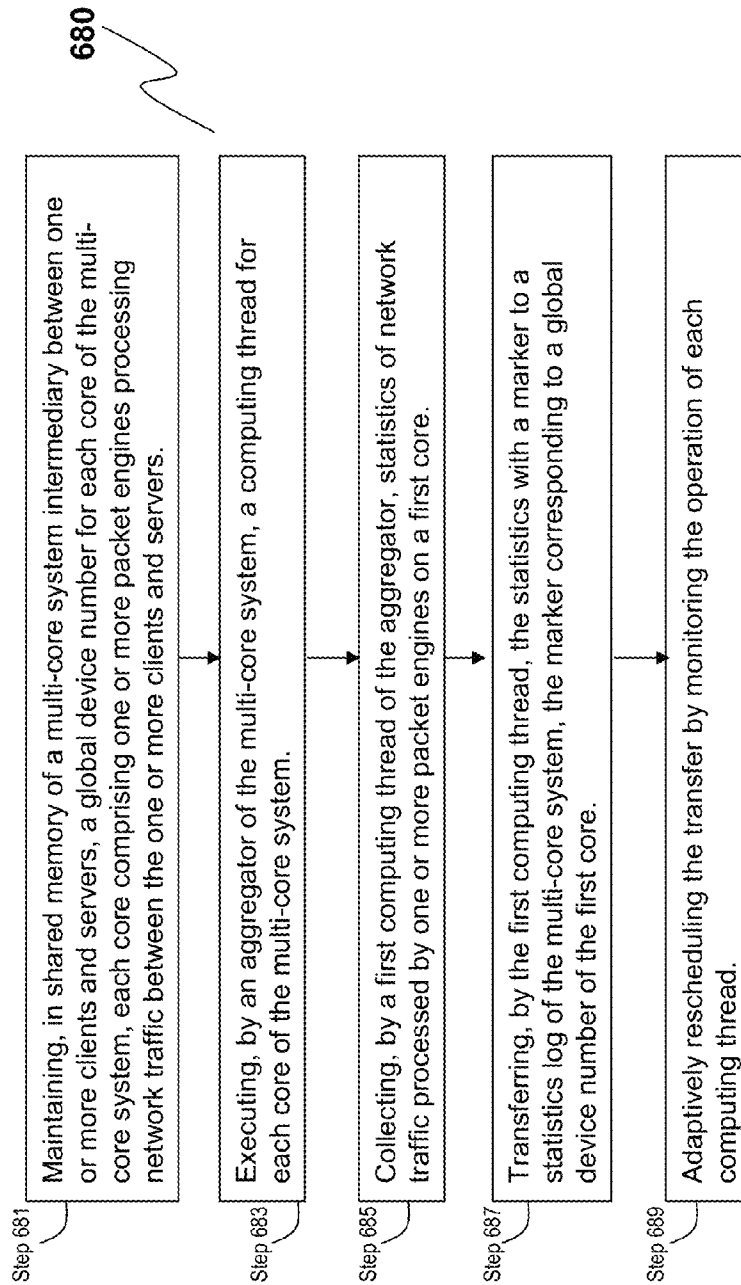


FIG. 6D

**FIG. 6E**

1

SYSTEMS AND METHODS FOR SCALABLE N-CORE STATS AGGREGATION

RELATED APPLICATIONS

The present application claims priority to U.S. Provisional Patent Application No. 61/428,124, entitled "SYSTEMS AND METHODS FOR SCALABLE N-CORE STATS AGGREGATION", filed Dec. 29, 2010, which is incorporated herein by reference in its entirety for all purposes.

FIELD OF THE DISCLOSURE

The present application generally relates to data communication networks. In particular, the present application relates to systems and methods for providing scalable unified data representation of the aggregate performance records and statistics of a multi-core network appliance.

BACKGROUND OF THE DISCLOSURE

In certain conventional network systems, an appliance or intermediary device is disposed in a network between one or more clients requesting web content or a network service, and one or more servers providing the requested web content or service. In some cases, the appliance establishes connections with the clients and servers and manages the connections and flow of information between the clients and servers. The appliance may employ security rules to ensure a measure of secure communications, monitor response time of servers, configure network connections to balance loads to servers, maintain user session data, and attend to other tasks which assure or improve the quality of communications supported by the appliance. In various implementations, an appliance or intermediary device may track and/or provide data about its operations or performance.

BRIEF SUMMARY OF THE DISCLOSURE

The present invention is directed towards systems and methods for aggregating and providing statistics from cores of a multi-core system intermediary between one or more clients and servers. The system may maintain, in shared memory, a global device number for each core of the multi-core system. The system may provide a thread for each core of the multi-core system to gather data from the corresponding core. A first thread may generate aggregated statistics from a corresponding core by parsing the gathered data from the corresponding core. The first thread may transfer the generated statistics to a statistics log. In some aspects, the present application is directed towards systems and methods for providing unified performance data and unified trace data for a multi-core appliance disposed in a network. In a multi-core appliance, a plurality of packet engines in operation on the appliance can manage network traffic between one or more clients and one or more servers. The systems and methods herein describe how the plurality of performance and trace data generated by the plurality of packet engines can be efficiently managed and consolidated to provide an aggregate set of performance data or trace data.

In one aspect, the present disclosure is directed to a method for aggregating performance statistics from multiple cores of a system intermediary between one or more clients and servers. The multiple cores may process multiple network traffic streams between one or more clients and servers. The method may include maintaining, in shared memory of a multi-core system intermediary between one or more clients and servers,

2

a global device number for each core of the multi-core system. Each core may include one or more packet engines processing network traffic between the one or more clients and servers. An aggregator of the multi-core system may execute a computing thread for each core of the multi-core system. A first computing thread of the aggregator may collect statistics of network traffic processed by one or more packet engines on a first core. The first computing thread may transfer the statistics with a marker to a statistics log of the multi-core system. The marker may correspond to a global device number of the first core.

In some embodiments, the aggregator may assign, for each core of the multi-core system, a global device number to a virtual machine executing on the core. The aggregator may assign, for each core of the multi-core system, a global device number to each virtual machine of the core, each virtual machine comprising one or more packet engines processing network traffic between the one or more clients and servers. The aggregator may consolidate at least a portion of statistics collected from two or more cores of the multi-core system into a buffer. The first computing thread may write the collected statistics to the statistics log according to a schedule of the aggregator. The aggregator may adaptively reschedule the transfer by monitoring the operation of the first core.

In some embodiments, the first computing thread may transfer the collected statistics unchanged to the statistics log. The multi-core system may maintain the statistics log for two or more cores of the multi-core system. The aggregator may send a notification to a first client of the one or more clients responsive to the transfer. The multi-core system may provide to the one or more clients access to the statistics log.

In another aspect, the present disclosure is directed to a system for aggregating performance statistics from multiple cores of a device intermediary between one or more clients and servers. The multiple cores may process multiple network traffic streams between one or more clients and servers. The system may include shared memory between multiple cores of the intermediary device, for maintaining a global device number for each core of the multi-core system. Each core may include one or more packet engines processing network traffic between the one or more clients and servers. An aggregator of the intermediary device may execute a computing thread for each core of the multi-core system. The aggregator may execute a first computing thread collecting statistics of network traffic processed by one or more packet engines on a first core. The aggregator may transfer the statistics with a marker to a statistics log of the multi-core system. The marker may correspond to a global device number of the first core.

In some embodiments, the aggregator assigns a global device number to a virtual machine executing on each core. The aggregator may assign a global device number to each virtual machine of each core. Each virtual machine may include one or more packet engines processing network traffic between the one or more clients and servers. The aggregator may further consolidate at least a portion of statistics collected from two or more cores of the multi-core system into a buffer. The first computing thread may write the collected statistics to the statistics log according to a schedule of the aggregator.

In some embodiments, the aggregator adaptively reschedules the transfer by monitoring the operation of the first core. The first computing thread may transfer the collected statistics unchanged to the statistics log. The aggregator may maintain the statistics log for two or more cores of the device. The aggregator may send a notification to a first client of the one

3

or more clients responsive to the transfer. The aggregator may provide the one or more clients access to the statistics log.

In various embodiments, a system for aggregating network performance data or trace data comprises an appliance disposed in a network between a plurality of clients and a plurality of servers providing web content or network services for clients, a plurality of packet engines in operation on a plurality of cores of the appliance, and an aggregator in operation on the appliance. In certain embodiments, a system for aggregating performance data of an appliance, which is intermediary to a plurality of clients and one or more servers in a networked system, comprises a plurality of packet engines executing on a plurality of cores configured for operation with the appliance. Each packet engine, executing on a respective core, can collect performance data identifying statistics of a service provided by a server in communication with the packet engine as well as data identifying statistics of a virtual service provided by the packet engine. One or more of the packet engines can manage network traffic associated with the service. In various embodiments, the system further comprises an aggregator which executes on a core of the multi-core, multi-packet-engine appliance. The aggregator can store to a buffer packet engine performance data obtained from each packet engine, and consolidate all packet engine performance data to identify unified performance data. In various embodiments, the unified performance data identifies statistics for network services managed by the appliance, in some embodiments as though the appliance comprised a single-core processing device. Further, the aggregator can receive a request from an agent for performance data of the appliance, and transmit the unified performance data in response to the request.

In certain embodiments, the system described above is adapted to aggregate trace data and provide a unified record of trace data for the multi-core, multi-packet-engine appliance. In such embodiments, each packet engine can capture trace information for network traffic received and/or transmitted by the packet engine. The aggregator can store to a buffer the trace information obtained from each packet engine, and consolidate all trace information from the plurality of packet engines to provide unified trace data of network traffic managed by the appliance. The unified trace data can be representative of the appliance as though the appliance were a single-core, single-packet-engine system. The aggregator can further receive a request from an agent for network trace data of the appliance, and transmit the unified trace data of network traffic in response to the request.

An embodiment of a method for aggregating performance data of a multi-core, multi-packet-engine appliance can comprise collecting, by each packet engine executing on a respective core of the appliance, packet engine performance data identifying statistics of a service provided by a server as managed by each packet engine, storing to a buffer, by the aggregator executing on a core of the appliance, the packet engine performance data obtained from each of the packet engines, and consolidating, by the aggregator, the packet engine performance data to identify unified performance data. The unified performance data can identify statistics for the service as managed by the appliance. The method can further comprise receiving, by the aggregator, a request from an agent for performance data of the appliance, and transmitting, by the aggregator, the unified performance data in response to the request.

In some embodiments, the step of collecting, by each packet engine, further comprises storing, by the aggregator, a first command in a first location of memory shared by the aggregator and a first core. The first command can identify a

4

first request for packet engine statistics from a first packet engine of the first core. The collecting can further comprise storing, by the aggregator, a second command in a second location of memory shared by the aggregator and a second core. The second command can identify a second request for packet engine statistics from a second packet engine of the second core. The step of storing, by the aggregator, can further comprise reading, by the aggregator, packet engine statistical data stored by the first packet engine in the memory shared between the aggregator and first core. The reading can be executed in response to detecting, by the aggregator, a change to a first flag in the shared memory. The storing can additionally comprise reading, by the aggregator, packet engine statistical data stored by the second packet engine in the memory shared between the aggregator and second core. The reading can be executed in response to detecting, by the aggregator, a change to a second flag in the shared memory.

In some embodiments, the step of collecting further comprises collecting, by each packet engine, local statistics of a virtual server of each packet engine. The virtual server can be configured to operate on the packet engine and manage the service provided by a server in communication with the packet engine. The step of storing can further comprise storing to a buffer, by the aggregator, the local statistics of each packet engine's virtual server obtained from each of the packet engines, and the step of consolidating of the method can further comprise consolidating, by the aggregator, each of the local statistics of each packet engine's virtual server to provide unified performance data identifying statistics for the appliance as a virtual server.

In some implementations, the step of collecting further comprises collecting, by each packet engine, a number of connections to the service of a server as packet engine performance data, and the step of consolidating, by the aggregator, further comprises consolidating, by the aggregator, the number of connections to the service from the plurality of packet engines. The aggregator can then provide a total number of connections to the service from the appliance. In some instances, the step of collecting further comprises collecting, by each packet engine, an average response time to the service as packet engine performance data, and the step of consolidating further comprises consolidating, by the aggregator, the average response to the service from the plurality of packet engines. The aggregator can provide a system-wide average response time to the service from the appliance. In some embodiments, the step of collecting further comprises collecting, by each packet engine, a number of bytes passed for the service as packet engine performance data, and the step of consolidating further comprises consolidating, by the aggregator, the number of bytes passed for the service from the plurality of packet engines. The aggregator can provide a total number of bytes passed for the service from the appliance. In certain embodiments, the step of collecting further comprises collecting, by each packet engine, a number of different servers providing service as packet engine performance data, and the step of consolidating further comprises consolidating the number of different servers from the plurality of packet engines. The aggregator can provide a total number of distinct servers providing service managed by the appliance.

An embodiment of a method for aggregating trace data of a multi-core, multi-packet-engine appliance can comprise capturing, by each packet engine executing on a respective core of the appliance, trace information for network traffic received or transmitted by the packet engine, storing to a buffer, by the aggregator, the trace information obtained from each packet engine of the plurality of packet engines, and consolidating, by the aggregator, the trace information from

5

each packet engine to provide unified trace data of network traffic managed by the appliance. The method for aggregating trace data can further comprise receiving, by the aggregator, a request from an agent for network trace data of the appliance, and transmitting, by the aggregator, the trace data of network traffic in response to the request. In various embodiments, the trace data may not identify the plurality of cores of the appliance. The agent requesting the data may not discern that the appliance comprises a plurality of cores and packet engines managing network traffic. In some embodiments, the method further comprises converting or structuring, by the aggregator, the trace data to a format recognizable by the agent requesting the data.

In some embodiments, the capturing, by each packet engine, further comprises storing, by the aggregator, a first command in a first location of memory shared with a first core. The first command can identify a first request for trace information from a first packet engine of the first core. The step of capturing can further comprise storing, by the aggregator, a second command in a first location of memory shared with a second core. The second command can identify a second request for trace information from a second packet engine of the second core. The step of storing, by the aggregator, can further comprise reading, by the aggregator, trace information stored by the first packet engine in the memory shared between the aggregator and the first core. The reading can be executed in response to detecting a change to a first flag in the shared memory monitored by the aggregator. The step of storing can further comprise reading, by the aggregator, trace information stored by the second packet engine in the memory shared between the aggregator and the second core. The reading can be executed in response to detecting a change to a second flag in the shared memory monitored by the aggregator.

In some embodiments, the step of capturing further comprises filtering, by each packet engine, trace information for network traffic, and the step of consolidating further comprises consolidating, by the aggregator, all the filtered trace information and providing a unified record of trace information for the appliance in accordance with the filtering. In some implementations, the step of capturing further comprises capturing, by each packet engine, trace information for network traffic from an identified source IP address, and the step of consolidating further comprises consolidating, by the aggregator, all the captured trace information and providing a unified record of network traffic from the identified source IP address managed by the appliance. In some instances of the method, the step of capturing further comprises capturing, by each packet engine, trace information for network traffic to an identified destination IP address, and the step of consolidating further comprises consolidating, by the aggregator, all the captured trace information and providing a unified record of network traffic to the identified destination IP address managed by the appliance.

The details of various embodiments of the invention are set forth in the accompanying drawings and the description below.

BRIEF DESCRIPTION OF THE FIGURES

The foregoing and other objects, aspects, features, and advantages of the invention will become more apparent and better understood by referring to the following description taken in conjunction with the accompanying drawings, in which:

FIG. 1A is a block diagram of an embodiment of a network environment for a client to access a server via an appliance;

6

FIG. 1B is a block diagram of an embodiment of an environment for delivering a computing environment from a server to a client via an appliance;

FIG. 1C is a block diagram of another embodiment of an environment for delivering a computing environment from a server to a client via an appliance;

FIG. 1D is a block diagram of another embodiment of an environment for delivering a computing environment from a server to a client via an appliance;

FIGS. 1E-1H are block diagrams of embodiments of a computing device;

FIG. 2A is a block diagram of an embodiment of an appliance for processing communications between a client and a server;

FIG. 2B is a block diagram of another embodiment of an appliance for optimizing, accelerating, load-balancing and routing communications between a client and a server;

FIG. 3 is a block diagram of an embodiment of a client for communicating with a server via the appliance;

FIG. 4A is a block diagram of an embodiment of a virtualization environment;

FIG. 4B is a block diagram of another embodiment of a virtualization environment;

FIG. 4C is a block diagram of an embodiment of a virtualized appliance;

FIG. 5A are block diagrams of embodiments of approaches to implementing parallelism in a multi-core system;

FIG. 5B is a block diagram of an embodiment of a system utilizing a multi-core system;

FIG. 5C is a block diagram of another embodiment of an aspect of a multi-core system;

FIG. 6A is a block diagram of an embodiment of a system for aggregating performance data and trace data for a multi-core, multi-packet-engine appliance;

FIG. 6B is an embodiment of a method for aggregating performance data or trace data for a multi-core, multi-packet-engine appliance;

FIG. 6C is an embodiment of additional steps for a method of aggregating performance data or trace data in a multi-core, multi-packet-engine appliance;

FIG. 6D is a block diagram of another embodiment of a system for aggregating performance data and trace data for a multi-core appliance; and

FIG. 6E is an embodiment of a method for aggregating performance data and/or trace data for a multi-core appliance.

The features and advantages of the present invention will become more apparent from the detailed description set forth below when taken in conjunction with the drawings, in which like reference characters identify corresponding elements throughout. In the drawings, like reference numbers generally indicate identical, functionally similar, and/or structurally similar elements.

DETAILED DESCRIPTION OF THE DISCLOSURE

For purposes of reading the description of the various embodiments below, the following descriptions of the sections of the specification and their respective contents may be helpful:

Section A describes a network environment and computing environment which may be useful for practicing embodiments described herein;

Section B describes embodiments of systems and methods for delivering a computing environment to a remote user;

Section C describes embodiments of systems and methods for accelerating communications between a client and a server;

Section D describes embodiments of systems and methods for virtualizing an application delivery controller;

Section E describes embodiments of systems and methods for providing a multi-core architecture and environment;

Section F describes embodiments of systems and methods for aggregating performance and trace data in a multi-core system; and

Section G describes systems and methods for scalable multi-core stats aggregation.

A. Network and Computing Environment

Prior to discussing the specifics of embodiments of the systems and methods of an appliance and/or client, it may be helpful to discuss the network and computing environments in which such embodiments may be deployed. Referring now to FIG. 1A, an embodiment of a network environment is depicted. In brief overview, the network environment comprises one or more clients **102a-102n** (also generally referred to as local machine(s) **102**, or client(s) **102**) in communication with one or more servers **106a-106n** (also generally referred to as server(s) **106**, or remote machine(s) **106**) via one or more networks **104**, **104'** (generally referred to as network **104**). In some embodiments, a client **102** communicates with a server **106** via an appliance **200**.

Although FIG. 1A shows a network **104** and a network **104'** between the clients **102** and the servers **106**, the clients **102** and the servers **106** may be on the same network **104**. The networks **104** and **104'** can be the same type of network or different types of networks. The network **104** and/or the network **104'** can be a local-area network (LAN), such as a company Intranet, a metropolitan area network (MAN), or a wide area network (WAN), such as the Internet or the World Wide Web. In one embodiment, network **104'** may be a private network and network **104** may be a public network. In some embodiments, network **104** may be a private network and network **104'** a public network. In another embodiment, networks **104** and **104'** may both be private networks. In some embodiments, clients **102** may be located at a branch office of a corporate enterprise communicating via a WAN connection over the network **104** to the servers **106** located at a corporate data center.

The network **104** and/or **104'** be any type and/or form of network and may include any of the following: a point to point network, a broadcast network, a wide area network, a local area network, a telecommunications network, a data communication network, a computer network, an ATM (Asynchronous Transfer Mode) network, a SONET (Synchronous Optical Network) network, a SDH (Synchronous Digital Hierarchy) network, a wireless network and a wireline network. In some embodiments, the network **104** may comprise a wireless link, such as an infrared channel or satellite band. The topology of the network **104** and/or **104'** may be a bus, star, or ring network topology. The network **104** and/or **104'** and network topology may be of any such network or network topology as known to those ordinarily skilled in the art capable of supporting the operations described herein.

As shown in FIG. 1A, the appliance **200**, which also may be referred to as an interface unit **200** or gateway **200**, is shown between the networks **104** and **104'**. In some embodiments, the appliance **200** may be located on network **104**. For example, a branch office of a corporate enterprise may deploy an appliance **200** at the branch office. In other embodiments, the appliance **200** may be located on network **104'**. For example, an appliance **200** may be located at a corporate data center. In yet another embodiment, a plurality of appliances

200 may be deployed on network **104**. In some embodiments, a plurality of appliances **200** may be deployed on network **104'**. In one embodiment, a first appliance **200** communicates with a second appliance **200'**. In other embodiments, the appliance **200** could be a part of any client **102** or server **106** on the same or different network **104,104'** as the client **102**. One or more appliances **200** may be located at any point in the network or network communications path between a client **102** and a server **106**.

In some embodiments, the appliance **200** comprises any of the network devices manufactured by Citrix Systems, Inc. of Ft. Lauderdale Fla., referred to as Citrix NetScaler devices. In other embodiments, the appliance **200** includes any of the product embodiments referred to as WebAccelerator and BigIP manufactured by F5 Networks, Inc. of Seattle, Wash. In another embodiment, the appliance **205** includes any of the DX acceleration device platforms and/or the SSL VPN series of devices, such as SA 700, SA 2000, SA 4000, and SA 6000 devices manufactured by Juniper Networks, Inc. of Sunnyvale, Calif. In yet another embodiment, the appliance **200** includes any application acceleration and/or security related appliances and/or software manufactured by Cisco Systems, Inc. of San Jose, Calif., such as the Cisco ACE Application Control Engine Module service software and network modules, and Cisco AVS Series Application Velocity System.

In one embodiment, the system may include multiple, logically-grouped servers **106**. In these embodiments, the logical group of servers may be referred to as a server farm **38**. In some of these embodiments, the servers **106** may be geographically dispersed. In some cases, a farm **38** may be administered as a single entity. In other embodiments, the server farm **38** comprises a plurality of server farms **38**. In one embodiment, the server farm executes one or more applications on behalf of one or more clients **102**.

The servers **106** within each farm **38** can be heterogeneous. One or more of the servers **106** can operate according to one type of operating system platform (e.g., WINDOWS NT, manufactured by Microsoft Corp. of Redmond, Wash.), while one or more of the other servers **106** can operate on according to another type of operating system platform (e.g., Unix or Linux). The servers **106** of each farm **38** do not need to be physically proximate to another server **106** in the same farm **38**. Thus, the group of servers **106** logically grouped as a farm **38** may be interconnected using a wide-area network (WAN) connection or medium-area network (MAN) connection. For example, a farm **38** may include servers **106** physically located in different continents or different regions of a continent, country, state, city, campus, or room. Data transmission speeds between servers **106** in the farm **38** can be increased if the servers **106** are connected using a local-area network (LAN) connection or some form of direct connection.

Servers **106** may be referred to as a file server, application server, web server, proxy server, or gateway server. In some embodiments, a server **106** may have the capacity to function as either an application server or as a master application server. In one embodiment, a server **106** may include an Active Directory. The clients **102** may also be referred to as client nodes or endpoints. In some embodiments, a client **102** has the capacity to function as both a client node seeking access to applications on a server and as an application server providing access to hosted applications for other clients **102a-102n**.

In some embodiments, a client **102** communicates with a server **106**. In one embodiment, the client **102** communicates directly with one of the servers **106** in a farm **38**. In another embodiment, the client **102** executes a program neighborhood application to communicate with a server **106** in a farm

38. In still another embodiment, the server 106 provides the functionality of a master node. In some embodiments, the client 102 communicates with the server 106 in the farm 38 through a network 104. Over the network 104, the client 102 can, for example, request execution of various applications hosted by the servers 106a-106n in the farm 38 and receive output of the results of the application execution for display. In some embodiments, only the master node provides the functionality required to identify and provide address information associated with a server 106' hosting a requested application.

In one embodiment, the server 106 provides functionality of a web server. In another embodiment, the server 106a receives requests from the client 102, forwards the requests to a second server 106b and responds to the request by the client 102 with a response to the request from the server 106b. In still another embodiment, the server 106 acquires an enumeration of applications available to the client 102 and address information associated with a server 106 hosting an application identified by the enumeration of applications. In yet another embodiment, the server 106 presents the response to the request to the client 102 using a web interface. In one embodiment, the client 102 communicates directly with the server 106 to access the identified application. In another embodiment, the client 102 receives application output data, such as display data, generated by an execution of the identified application on the server 106.

Referring now to FIG. 1B, an embodiment of a network environment deploying multiple appliances 200 is depicted. A first appliance 200 may be deployed on a first network 104 and a second appliance 200' on a second network 104'. For example a corporate enterprise may deploy a first appliance 200 at a branch office and a second appliance 200' at a data center. In another embodiment, the first appliance 200 and second appliance 200' are deployed on the same network 104 or network 104'. For example, a first appliance 200 may be deployed for a first server farm 38, and a second appliance 200 may be deployed for a second server farm 38'. In another example, a first appliance 200 may be deployed at a first branch office while the second appliance 200' is deployed at a second branch office'. In some embodiments, the first appliance 200 and second appliance 200' work in cooperation or in conjunction with each other to accelerate network traffic or the delivery of application and data between a client and a server.

Referring now to FIG. 1C, another embodiment of a network environment deploying the appliance 200 with one or more other types of appliances, such as between one or more WAN optimization appliance 205, 205' is depicted. For example a first WAN optimization appliance 205 is shown between networks 104 and 104' and a second WAN optimization appliance 205' may be deployed between the appliance 200 and one or more servers 106. By way of example, a corporate enterprise may deploy a first WAN optimization appliance 205 at a branch office and a second WAN optimization appliance 205' at a data center. In some embodiments, the appliance 205 may be located on network 104'. In other embodiments, the appliance 205' may be located on network 104. In some embodiments, the appliance 205' may be located on network 104' or network 104". In one embodiment, the appliance 205 and 205' are on the same network. In another embodiment, the appliance 205 and 205' are on different networks. In another example, a first WAN optimization appliance 205 may be deployed for a first server farm 38 and a second WAN optimization appliance 205' for a second server farm 38'.

In one embodiment, the appliance 205 is a device for accelerating, optimizing or otherwise improving the performance, operation, or quality of service of any type and form of network traffic, such as traffic to and/or from a WAN connection. In some embodiments, the appliance 205 is a performance enhancing proxy. In other embodiments, the appliance 205 is any type and form of WAN optimization or acceleration device, sometimes also referred to as a WAN optimization controller. In one embodiment, the appliance 205 is any of the product embodiments referred to as WANScaler manufactured by Citrix Systems, Inc. of Ft. Lauderdale, Fla. In other embodiments, the appliance 205 includes any of the product embodiments referred to as BIG-IP link controller and WANjet manufactured by F5 Networks, Inc. of Seattle, Wash. In another embodiment, the appliance 205 includes any of the WX and WXC WAN acceleration device platforms manufactured by Juniper Networks, Inc. of Sunnyvale, Calif. In some embodiments, the appliance 205 includes any of the steelhead line of WAN optimization appliances manufactured by Riverbed Technology of San Francisco, Calif. In other embodiments, the appliance 205 includes any of the WAN related devices manufactured by Expand Networks Inc. of Roseland, N.J. In one embodiment, the appliance 205 includes any of the WAN related appliances manufactured by Packeteer Inc. of Cupertino, Calif., such as the PacketShaper, iShared, and SkyX product embodiments provided by Packeteer. In yet another embodiment, the appliance 205 includes any WAN related appliances and/or software manufactured by Cisco Systems, Inc. of San Jose, Calif., such as the Cisco Wide Area Network Application Services software and network modules, and Wide Area Network engine appliances.

In one embodiment, the appliance 205 provides application and data acceleration services for branch-office or remote offices. In one embodiment, the appliance 205 includes optimization of Wide Area File Services (WAFS). In another embodiment, the appliance 205 accelerates the delivery of files, such as via the Common Internet File System (CIFS) protocol. In other embodiments, the appliance 205 provides caching in memory and/or storage to accelerate delivery of applications and data. In one embodiment, the appliance 205 provides compression of network traffic at any level of the network stack or at any protocol or network layer. In another embodiment, the appliance 205 provides transport layer protocol optimizations, flow control, performance enhancements or modifications and/or management to accelerate delivery of applications and data over a WAN connection. For example, in one embodiment, the appliance 205 provides Transport Control Protocol (TCP) optimizations. In other embodiments, the appliance 205 provides optimizations, flow control, performance enhancements or modifications and/or management for any session or application layer protocol.

In another embodiment, the appliance 205 encoded any type and form of data or information into custom or standard TCP and/or IP header fields or option fields of network packet to announce presence, functionality or capability to another appliance 205'. In another embodiment, an appliance 205' may communicate with another appliance 205' using data encoded in both TCP and/or IP header fields or options. For example, the appliance may use TCP option(s) or IP header fields or options to communicate one or more parameters to be used by the appliances 205, 205' in performing functionality, such as WAN acceleration, or for working in conjunction with each other.

In some embodiments, the appliance 200 preserves any of the information encoded in TCP and/or IP header and/or option fields communicated between appliances 205 and 205'. For example, the appliance 200 may terminate a trans-

11

port layer connection traversing the appliance **200**, such as a transport layer connection from between a client and a server traversing appliances **205** and **205'**. In one embodiment, the appliance **200** identifies and preserves any encoded information in a transport layer packet transmitted by a first appliance **205** via a first transport layer connection and communicates a transport layer packet with the encoded information to a second appliance **205'** via a second transport layer connection.

Referring now to FIG. 1D, a network environment for delivering and/or operating a computing environment on a client **102** is depicted. In some embodiments, a server **106** includes an application delivery system **190** for delivering a computing environment or an application and/or data file to one or more clients **102**. In brief overview, a client **102** is in communication with a server **106** via network **104**, **104'** and appliance **200**. For example, the client **102** may reside in a remote office of a company, e.g., a branch office, and the server **106** may reside at a corporate data center. The client **102** comprises a client agent **120**, and a computing environment **15**. The computing environment **15** may execute or operate an application that accesses, processes or uses a data file. The computing environment **15**, application and/or data file may be delivered via the appliance **200** and/or the server **106**.

In some embodiments, the appliance **200** accelerates delivery of a computing environment **15**, or any portion thereof, to a client **102**. In one embodiment, the appliance **200** accelerates the delivery of the computing environment **15** by the application delivery system **190**. For example, the embodiments described herein may be used to accelerate delivery of a streaming application and data file processable by the application from a central corporate data center to a remote user location, such as a branch office of the company. In another embodiment, the appliance **200** accelerates transport layer traffic between a client **102** and a server **106**. The appliance **200** may provide acceleration techniques for accelerating any transport layer payload from a server **106** to a client **102**, such as: 1) transport layer connection pooling, 2) transport layer connection multiplexing, 3) transport control protocol buffering, 4) compression and 5) caching. In some embodiments, the appliance **200** provides load balancing of servers **106** in responding to requests from clients **102**. In other embodiments, the appliance **200** acts as a proxy or access server to provide access to the one or more servers **106**. In another embodiment, the appliance **200** provides a secure virtual private network connection from a first network **104** of the client **102** to the second network **104'** of the server **106**, such as an SSL VPN connection. In yet other embodiments, the appliance **200** provides application firewall security, control and management of the connection and communications between a client **102** and a server **106**.

In some embodiments, the application delivery management system **190** provides application delivery techniques to deliver a computing environment to a desktop of a user, remote or otherwise, based on a plurality of execution methods and based on any authentication and authorization policies applied via a policy engine **195**. With these techniques, a remote user may obtain a computing environment and access to server stored applications and data files from any network connected device **100**. In one embodiment, the application delivery system **190** may reside or execute on a server **106**. In another embodiment, the application delivery system **190** may reside or execute on a plurality of servers **106a-106n**. In some embodiments, the application delivery system **190** may execute in a server farm **38**. In one embodiment, the server **106** executing the application delivery system **190** may also

12

store or provide the application and data file. In another embodiment, a first set of one or more servers **106** may execute the application delivery system **190**, and a different server **106n** may store or provide the application and data file. In some embodiments, each of the application delivery system **190**, the application, and data file may reside or be located on different servers. In yet another embodiment, any portion of the application delivery system **190** may reside, execute or be stored on or distributed to the appliance **200**, or a plurality of appliances.

The client **102** may include a computing environment **15** for executing an application that uses or processes a data file. The client **102** via networks **104**, **104'** and appliance **200** may request an application and data file from the server **106**. In one embodiment, the appliance **200** may forward a request from the client **102** to the server **106**. For example, the client **102** may not have the application and data file stored or accessible locally. In response to the request, the application delivery system **190** and/or server **106** may deliver the application and data file to the client **102**. For example, in one embodiment, the server **106** may transmit the application as an application stream to operate in computing environment **15** on client **102**.

In some embodiments, the application delivery system **190** comprises any portion of the Citrix Access Suite™ by Citrix Systems, Inc., such as the MetaFrame or Citrix Presentation Server™ and/or any of the Microsoft® Windows Terminal Services manufactured by the Microsoft Corporation. In one embodiment, the application delivery system **190** may deliver one or more applications to clients **102** or users via a remote-display protocol or otherwise via remote-based or server-based computing. In another embodiment, the application delivery system **190** may deliver one or more applications to clients or users via streaming of the application.

In one embodiment, the application delivery system **190** includes a policy engine **195** for controlling and managing the access to, selection of application execution methods and the delivery of applications. In some embodiments, the policy engine **195** determines the one or more applications a user or client **102** may access. In another embodiment, the policy engine **195** determines how the application should be delivered to the user or client **102**, e.g., the method of execution. In some embodiments, the application delivery system **190** provides a plurality of delivery techniques from which to select a method of application execution, such as a server-based computing, streaming or delivering the application locally to the client **120** for local execution.

In one embodiment, a client **102** requests execution of an application program and the application delivery system **190** comprising a server **106** selects a method of executing the application program. In some embodiments, the server **106** receives credentials from the client **102**. In another embodiment, the server **106** receives a request for an enumeration of available applications from the client **102**. In one embodiment, in response to the request or receipt of credentials, the application delivery system **190** enumerates a plurality of application programs available to the client **102**. The application delivery system **190** receives a request to execute an enumerated application. The application delivery system **190** selects one of a predetermined number of methods for executing the enumerated application, for example, responsive to a policy of a policy engine. The application delivery system **190** may select a method of execution of the application enabling the client **102** to receive application-output data generated by execution of the application program on a server **106**. The application delivery system **190** may select a method of execution of the application enabling the local machine **10** to execute the application program locally after

13

retrieving a plurality of application files comprising the application. In yet another embodiment, the application delivery system **190** may select a method of execution of the application to stream the application via the network **104** to the client **102**.

A client **102** may execute, operate or otherwise provide an application, which can be any type and/or form of software, program, or executable instructions such as any type and/or form of web browser, web-based client, client-server application, a thin-client computing client, an ActiveX control, or a Java applet, or any other type and/or form of executable instructions capable of executing on client **102**. In some embodiments, the application may be a server-based or a remote-based application executed on behalf of the client **102** on a server **106**. In one embodiment the server **106** may display output to the client **102** using any thin-client or remote-display protocol, such as the Independent Computing Architecture (ICA) protocol manufactured by Citrix Systems, Inc. of Ft. Lauderdale, Fla. or the Remote Desktop Protocol (RDP) manufactured by the Microsoft Corporation of Redmond, Wash. The application can use any type of protocol and it can be, for example, an HTTP client, an FTP client, an Oscar client, or a Telnet client. In other embodiments, the application comprises any type of software related to VoIP communications, such as a soft IP telephone. In further embodiments, the application comprises any application related to real-time data communications, such as applications for streaming video and/or audio.

In some embodiments, the server **106** or a server farm **38** may be running one or more applications, such as an application providing a thin-client computing or remote display presentation application. In one embodiment, the server **106** or server farm **38** executes as an application, any portion of the Citrix Access Suite™ by Citrix Systems, Inc., such as the MetaFrame or Citrix Presentation Server™, and/or any of the Microsoft® Windows Terminal Services manufactured by the Microsoft Corporation. In one embodiment, the application is an ICA client, developed by Citrix Systems, Inc. of Fort Lauderdale, Fla. In other embodiments, the application includes a Remote Desktop (RDP) client, developed by Microsoft Corporation of Redmond, Wash. Also, the server **106** may run an application, which for example, may be an application server providing email services such as Microsoft Exchange manufactured by the Microsoft Corporation of Redmond, Wash., a web or Internet server, or a desktop sharing server, or a collaboration server. In some embodiments, any of the applications may comprise any type of hosted service or products, such as GoToMeeting™ provided by Citrix Online Division, Inc. of Santa Barbara, Calif., WebEx™ provided by WebEx, Inc. of Santa Clara, Calif., or Microsoft Office Live Meeting provided by Microsoft Corporation of Redmond, Wash.

Still referring to FIG. 1D, an embodiment of the network environment may include a monitoring server **106A**. The monitoring server **106A** may include any type and form performance monitoring service **198**. The performance monitoring service **198** may include monitoring, measurement and/or management software and/or hardware, including data collection, aggregation, analysis, management and reporting. In one embodiment, the performance monitoring service **198** includes one or more monitoring agents **197**. The monitoring agent **197** includes any software, hardware or combination thereof for performing monitoring, measurement and data collection activities on a device, such as a client **102**, server **106** or an appliance **200**, **205**. In some embodiments, the monitoring agent **197** includes any type and form of script, such as Visual Basic script, or Javascript. In one embodiment,

14

the monitoring agent **197** executes transparently to any application and/or user of the device. In some embodiments, the monitoring agent **197** is installed and operated unobtrusively to the application or client. In yet another embodiment, the monitoring agent **197** is installed and operated without any instrumentation for the application or device.

In some embodiments, the monitoring agent **197** monitors, measures and collects data on a predetermined frequency. In other embodiments, the monitoring agent **197** monitors, measures and collects data based upon detection of any type and form of event. For example, the monitoring agent **197** may collect data upon detection of a request for a web page or receipt of an HTTP response. In another example, the monitoring agent **197** may collect data upon detection of any user input events, such as a mouse click. The monitoring agent **197** may report or provide any monitored, measured or collected data to the monitoring service **198**. In one embodiment, the monitoring agent **197** transmits information to the monitoring service **198** according to a schedule or a predetermined frequency. In another embodiment, the monitoring agent **197** transmits information to the monitoring service **198** upon detection of an event.

In some embodiments, the monitoring service **198** and/or monitoring agent **197** performs monitoring and performance measurement of any network resource or network infrastructure element, such as a client, server, server farm, appliance **200**, appliance **205**, or network connection. In one embodiment, the monitoring service **198** and/or monitoring agent **197** performs monitoring and performance measurement of any transport layer connection, such as a TCP or UDP connection. In another embodiment, the monitoring service **198** and/or monitoring agent **197** monitors and measures network latency. In yet one embodiment, the monitoring service **198** and/or monitoring agent **197** monitors and measures bandwidth utilization.

In other embodiments, the monitoring service **198** and/or monitoring agent **197** monitors and measures end-user response times. In some embodiments, the monitoring service **198** performs monitoring and performance measurement of an application. In another embodiment, the monitoring service **198** and/or monitoring agent **197** performs monitoring and performance measurement of any session or connection to the application. In one embodiment, the monitoring service **198** and/or monitoring agent **197** monitors and measures performance of a browser. In another embodiment, the monitoring service **198** and/or monitoring agent **197** monitors and measures performance of HTTP based transactions. In some embodiments, the monitoring service **198** and/or monitoring agent **197** monitors and measures performance of a Voice over IP (VoIP) application or session. In other embodiments, the monitoring service **198** and/or monitoring agent **197** monitors and measures performance of a remote display protocol application, such as an ICA client or RDP client. In yet another embodiment, the monitoring service **198** and/or monitoring agent **197** monitors and measures performance of any type and form of streaming media. In still a further embodiment, the monitoring service **198** and/or monitoring agent **197** monitors and measures performance of a hosted application or a Software-As-A-Service (SaaS) delivery model.

In some embodiments, the monitoring service **198** and/or monitoring agent **197** performs monitoring and performance measurement of one or more transactions, requests or responses related to application. In other embodiments, the monitoring service **198** and/or monitoring agent **197** monitors and measures any portion of an application layer stack, such as any .NET or J2EE calls. In one embodiment, the

15

monitoring service **198** and/or monitoring agent **197** monitors and measures database or SQL transactions. In yet another embodiment, the monitoring service **198** and/or monitoring agent **197** monitors and measures any method, function or application programming interface (API) call.

In one embodiment, the monitoring service **198** and/or monitoring agent **197** performs monitoring and performance measurement of a delivery of application and/or data from a server to a client via one or more appliances, such as appliance **200** and/or appliance **205**. In some embodiments, the monitoring service **198** and/or monitoring agent **197** monitors and measures performance of delivery of a virtualized application. In other embodiments, the monitoring service **198** and/or monitoring agent **197** monitors and measures performance of delivery of a streaming application. In another embodiment, the monitoring service **198** and/or monitoring agent **197** monitors and measures performance of delivery of a desktop application to a client and/or the execution of the desktop application on the client. In another embodiment, the monitoring service **198** and/or monitoring agent **197** monitors and measures performance of a client/server application.

In one embodiment, the monitoring service **198** and/or monitoring agent **197** is designed and constructed to provide application performance management for the application delivery system **190**. For example, the monitoring service **198** and/or monitoring agent **197** may monitor, measure and manage the performance of the delivery of applications via the Citrix Presentation Server. In this example, the monitoring service **198** and/or monitoring agent **197** monitors individual ICA sessions. The monitoring service **198** and/or monitoring agent **197** may measure the total and per session system resource usage, as well as application and networking performance. The monitoring service **198** and/or monitoring agent **197** may identify the active servers for a given user and/or user session. In some embodiments, the monitoring service **198** and/or monitoring agent **197** monitors back-end connections between the application delivery system **190** and an application and/or database server. The monitoring service **198** and/or monitoring agent **197** may measure network latency, delay and volume per user-session or ICA session.

In some embodiments, the monitoring service **198** and/or monitoring agent **197** measures and monitors memory usage for the application delivery system **190**, such as total memory usage, per user session and/or per process. In other embodiments, the monitoring service **198** and/or monitoring agent **197** measures and monitors CPU usage the application delivery system **190**, such as total CPU usage, per user session and/or per process. In another embodiment, the monitoring service **198** and/or monitoring agent **197** measures and monitors the time required to log-in to an application, a server, or the application delivery system, such as Citrix Presentation Server. In one embodiment, the monitoring service **198** and/or monitoring agent **197** measures and monitors the duration a user is logged into an application, a server, or the application delivery system **190**. In some embodiments, the monitoring service **198** and/or monitoring agent **197** measures and monitors active and inactive session counts for an application, server or application delivery system session. In yet another embodiment, the monitoring service **198** and/or monitoring agent **197** measures and monitors user session latency.

In yet further embodiments, the monitoring service **198** and/or monitoring agent **197** measures and monitors measures and monitors any type and form of server metrics. In one embodiment, the monitoring service **198** and/or monitoring agent **197** measures and monitors metrics related to system memory, CPU usage, and disk storage. In another embodi-

16

ment, the monitoring service **198** and/or monitoring agent **197** measures and monitors metrics related to page faults, such as page faults per second. In other embodiments, the monitoring service **198** and/or monitoring agent **197** measures and monitors round-trip time metrics. In yet another embodiment, the monitoring service **198** and/or monitoring agent **197** measures and monitors metrics related to application crashes, errors and/or hangs.

In some embodiments, the monitoring service **198** and monitoring agent **198** includes any of the product embodiments referred to as EdgeSight manufactured by Citrix Systems, Inc. of Ft. Lauderdale, Fla. In another embodiment, the performance monitoring service **198** and/or monitoring agent **198** includes any portion of the product embodiments referred to as the TrueView product suite manufactured by the Symphoniq Corporation of Palo Alto, Calif. In one embodiment, the performance monitoring service **198** and/or monitoring agent **198** includes any portion of the product embodiments referred to as the TeaLeaf CX product suite manufactured by the TeaLeaf Technology Inc. of San Francisco, Calif. In other embodiments, the performance monitoring service **198** and/or monitoring agent **198** includes any portion of the business service management products, such as the BMC Performance Manager and Patrol products, manufactured by BMC Software, Inc. of Houston, Tex.

The client **102**, server **106**, and appliance **200** may be deployed as and/or executed on any type and form of computing device, such as a computer, network device or appliance capable of communicating on any type and form of network and performing the operations described herein. FIGS. 1E and 1F depict block diagrams of a computing device **100** useful for practicing an embodiment of the client **102**, server **106** or appliance **200**. As shown in FIGS. 1E and 1F, each computing device **100** includes a central processing unit **101**, and a main memory unit **122**. As shown in FIG. 1E, a computing device **100** may include a visual display device **124**, a keyboard **126** and/or a pointing device **127**, such as a mouse. Each computing device **100** may also include additional optional elements, such as one or more input/output devices **130a-130b** (generally referred to using reference numeral **130**), and a cache memory **140** in communication with the central processing unit **101**.

The central processing unit **101** is any logic circuitry that responds to and processes instructions fetched from the main memory unit **122**. In many embodiments, the central processing unit is provided by a microprocessor unit, such as: those manufactured by Intel Corporation of Mountain View, Calif.; those manufactured by Motorola Corporation of Schaumburg, Ill.; those manufactured by Transmeta Corporation of Santa Clara, Calif.; the RS/6000 processor, those manufactured by International Business Machines of White Plains, N.Y.; or those manufactured by Advanced Micro Devices of Sunnyvale, Calif. The computing device **100** may be based on any of these processors, or any other processor capable of operating as described herein.

Main memory unit **122** may be one or more memory chips capable of storing data and allowing any storage location to be directly accessed by the microprocessor **101**, such as Static random access memory (SRAM), Burst SRAM or Synchronous Burst SRAM (BSRAM), Dynamic random access memory (DRAM), Fast Page Mode DRAM (FPM DRAM), Enhanced DRAM (EDRAM), Extended Data Output RAM (EDO RAM), Extended Data Output DRAM (EDO DRAM), Burst Extended Data Output DRAM (BEDO DRAM), Enhanced DRAM (EDRAM), synchronous DRAM (SDRAM), JEDEC SRAM, PC100 SDRAM, Double Data Rate SDRAM (DDR SDRAM), Enhanced SDRAM (ESDRAM), SyncLink

17

DRAM (SLDRAM), Direct Rambus DRAM (DRDRAM), or Ferroelectric RAM (FRAM). The main memory **122** may be based on any of the above described memory chips, or any other available memory chips capable of operating as described herein. In the embodiment shown in FIG. 1E, the processor **101** communicates with main memory **122** via a system bus **150** (described in more detail below). FIG. 1F depicts an embodiment of a computing device **100** in which the processor communicates directly with main memory **122** via a memory port **103**. For example, in FIG. 1F the main memory **122** may be DRDRAM.

FIG. 1F depicts an embodiment in which the main processor **101** communicates directly with cache memory **140** via a secondary bus, sometimes referred to as a backside bus. In other embodiments, the main processor **101** communicates with cache memory **140** using the system bus **150**. Cache memory **140** typically has a faster response time than main memory **122** and is typically provided by SRAM, BSRAM, or EDRAM. In the embodiment shown in FIG. 1F, the processor **101** communicates with various I/O devices **130** via a local system bus **150**. Various busses may be used to connect the central processing unit **101** to any of the I/O devices **130**, including a VESA VL bus, an ISA bus, an EISA bus, a MicroChannel Architecture (MCA) bus, a PCI bus, a PCI-X bus, a PCI-Express bus, or a NuBus. For embodiments in which the I/O device is a video display **124**, the processor **101** may use an Advanced Graphics Port (AGP) to communicate with the display **124**. FIG. 1F depicts an embodiment of a computer **100** in which the main processor **101** communicates directly with I/O device **130b** via HyperTransport, Rapid I/O, or InfiniBand. FIG. 1F also depicts an embodiment in which local busses and direct communication are mixed: the processor **101** communicates with I/O device **130b** using a local interconnect bus while communicating with I/O device **130a** directly.

The computing device **100** may support any suitable installation device **116**, such as a floppy disk drive for receiving floppy disks such as 3.5-inch, 5.25-inch disks or ZIP disks, a CD-ROM drive, a CD-R/RW drive, a DVD-ROM drive, tape drives of various formats, USB device, hard-drive or any other device suitable for installing software and programs such as any client agent **120**, or portion thereof. The computing device **100** may further comprise a storage device **128**, such as one or more hard disk drives or redundant arrays of independent disks, for storing an operating system and other related software, and for storing application software programs such as any program related to the client agent **120**. Optionally, any of the installation devices **116** could also be used as the storage device **128**. Additionally, the operating system and the software can be run from a bootable medium, for example, a bootable CD, such as KNOPPIX®, a bootable CD for GNU/Linux that is available as a GNU/Linux distribution from knoppix.net.

Furthermore, the computing device **100** may include a network interface **118** to interface to a Local Area Network (LAN), Wide Area Network (WAN) or the Internet through a variety of connections including, but not limited to, standard telephone lines, LAN or WAN links (e.g., 802.11, T1, T3, 56 kb, X.25), broadband connections (e.g., ISDN, Frame Relay, ATM), wireless connections, or some combination of any or all of the above. The network interface **118** may comprise a built-in network adapter, network interface card, PCMCIA network card, card bus network adapter, wireless network adapter, USB network adapter, modem or any other device suitable for interfacing the computing device **100** to any type of network capable of communication and performing the operations described herein. A wide variety of I/O devices

18

130a-130n may be present in the computing device **100**. Input devices include keyboards, mice, trackpads, trackballs, microphones, and drawing tablets. Output devices include video displays, speakers, inkjet printers, laser printers, and dye-sublimation printers. The I/O devices **130** may be controlled by an I/O controller **123** as shown in FIG. 1E. The I/O controller may control one or more I/O devices such as a keyboard **126** and a pointing device **127**, e.g., a mouse or optical pen. Furthermore, an I/O device may also provide storage **128** and/or an installation medium **116** for the computing device **100**. In still other embodiments, the computing device **100** may provide USB connections to receive handheld USB storage devices such as the USB Flash Drive line of devices manufactured by Twintech Industry, Inc. of Los Alamitos, Calif.

In some embodiments, the computing device **100** may comprise or be connected to multiple display devices **124a-124n**, which each may be of the same or different type and/or form. As such, any of the I/O devices **130a-130n** and/or the I/O controller **123** may comprise any type and/or form of suitable hardware, software, or combination of hardware and software to support, enable or provide for the connection and use of multiple display devices **124a-124n** by the computing device **100**. For example, the computing device **100** may include any type and/or form of video adapter, video card, driver, and/or library to interface, communicate, connect or otherwise use the display devices **124a-124n**. In one embodiment, a video adapter may comprise multiple connectors to interface to multiple display devices **124a-124n**. In other embodiments, the computing device **100** may include multiple video adapters, with each video adapter connected to one or more of the display devices **124a-124n**. In some embodiments, any portion of the operating system of the computing device **100** may be configured for using multiple displays **124a-124n**. In other embodiments, one or more of the display devices **124a-124n** may be provided by one or more other computing devices, such as computing devices **100a** and **100b** connected to the computing device **100**, for example, via a network. These embodiments may include any type of software designed and constructed to use another computer's display device as a second display device **124a** for the computing device **100**. One ordinarily skilled in the art will recognize and appreciate the various ways and embodiments that a computing device **100** may be configured to have multiple display devices **124a-124n**.

In further embodiments, an I/O device **130** may be a bridge **170** between the system bus **150** and an external communication bus, such as a USB bus, an Apple Desktop Bus, an RS-232 serial connection, a SCSI bus, a FireWire bus, a FireWire **800** bus, an Ethernet bus, an AppleTalk bus, a Gigabit Ethernet bus, an Asynchronous Transfer Mode bus, a HIPPI bus, a Super HIPPI bus, a SerialPlus bus, a SCI/LAMP bus, a FibreChannel bus, or a Serial Attached small computer system interface bus.

A computing device **100** of the sort depicted in FIGS. 1E and 1F typically operate under the control of operating systems, which control scheduling of tasks and access to system resources. The computing device **100** can be running any operating system such as any of the versions of the Microsoft® Windows operating systems, the different releases of the Unix and Linux operating systems, any version of the Mac OS® for Macintosh computers, any embedded operating system, any real-time operating system, any open source operating system, any proprietary operating system, any operating systems for mobile computing devices, or any other operating system capable of running on the computing device and performing the operations described herein. Typi-

cal operating systems include: WINDOWS 3.x, WINDOWS 95, WINDOWS 98, WINDOWS 2000, WINDOWS NT 3.51, WINDOWS NT 4.0, WINDOWS CE, and WINDOWS XP, all of which are manufactured by Microsoft Corporation of Redmond, Wash.; MacOS, manufactured by Apple Computer of Cupertino, California; OS/2, manufactured by International Business Machines of Armonk, N.Y.; and Linux, a freely-available operating system distributed by Caldera Corp. of Salt Lake City, Utah, or any type and/or form of a Unix operating system, among others.

In other embodiments, the computing device **100** may have different processors, operating systems, and input devices consistent with the device. For example, in one embodiment the computer **100** is a Treo 180, 270, 1060, 600 or 650 smart phone manufactured by Palm, Inc. In this embodiment, the Treo smart phone is operated under the control of the PalmOS operating system and includes a stylus input device as well as a five-way navigator device. Moreover, the computing device **100** can be any workstation, desktop computer, laptop or notebook computer, server, handheld computer, mobile telephone, any other computer, or other form of computing or telecommunications device that is capable of communication and that has sufficient processor power and memory capacity to perform the operations described herein.

As shown in FIG. 1G, the computing device **100** may comprise multiple processors and may provide functionality for simultaneous execution of instructions or for simultaneous execution of one instruction on more than one piece of data. In some embodiments, the computing device **100** may comprise a parallel processor with one or more cores. In one of these embodiments, the computing device **100** is a shared memory parallel device, with multiple processors and/or multiple processor cores, accessing all available memory as a single global address space. In another of these embodiments, the computing device **100** is a distributed memory parallel device with multiple processors each accessing local memory only. In still another of these embodiments, the computing device **100** has both some memory which is shared and some memory which can only be accessed by particular processors or subsets of processors. In still even another of these embodiments, the computing device **100**, such as a multi-core micro-processor, combines two or more independent processors into a single package, often a single integrated circuit (IC). In yet another of these embodiments, the computing device **100** includes a chip having a CELL BROADBAND ENGINE architecture and including a Power processor element and a plurality of synergistic processing elements, the Power processor element and the plurality of synergistic processing elements linked together by an internal high speed bus, which may be referred to as an element interconnect bus.

In some embodiments, the processors provide functionality for execution of a single instruction simultaneously on multiple pieces of data (SIMD). In other embodiments, the processors provide functionality for execution of multiple instructions simultaneously on multiple pieces of data (MIMD). In still other embodiments, the processor may use any combination of SIMD and MIMD cores in a single device.

In some embodiments, the computing device **100** may comprise a graphics processing unit. In one of these embodiments, depicted in FIG. 1H, the computing device **100** includes at least one central processing unit **101** and at least one graphics processing unit. In another of these embodiments, the computing device **100** includes at least one parallel processing unit and at least one graphics processing unit. In still another of these embodiments, the computing device **100**

includes a plurality of processing units of any type, one of the plurality of processing units comprising a graphics processing unit.

In some embodiments, a first computing device **100a** executes an application on behalf of a user of a client computing device **100b**. In other embodiments, a computing device **100a** executes a virtual machine, which provides an execution session within which applications execute on behalf of a user or a client computing devices **100b**. In one of these embodiments, the execution session is a hosted desktop session. In another of these embodiments, the computing device **100** executes a terminal services session. The terminal services session may provide a hosted desktop environment. In still another of these embodiments, the execution session provides access to a computing environment, which may comprise one or more of: an application, a plurality of applications, a desktop application, and a desktop session in which one or more applications may execute.

B. Appliance Architecture

FIG. 2A illustrates an example embodiment of the appliance **200**. The architecture of the appliance **200** in FIG. 2A is provided by way of illustration only and is not intended to be limiting. As shown in FIG. 2, appliance **200** comprises a hardware layer **206** and a software layer divided into a user space **202** and a kernel space **204**.

Hardware layer **206** provides the hardware elements upon which programs and services within kernel space **204** and user space **202** are executed. Hardware layer **206** also provides the structures and elements which allow programs and services within kernel space **204** and user space **202** to communicate data both internally and externally with respect to appliance **200**. As shown in FIG. 2, the hardware layer **206** includes a processing unit **262** for executing software programs and services, a memory **264** for storing software and data, network ports **266** for transmitting and receiving data over a network, and an encryption processor **260** for performing functions related to Secure Sockets Layer processing of data transmitted and received over the network. In some embodiments, the central processing unit **262** may perform the functions of the encryption processor **260** in a single processor. Additionally, the hardware layer **206** may comprise multiple processors for each of the processing unit **262** and the encryption processor **260**. The processor **262** may include any of the processors **101** described above in connection with FIGS. 1E and 1F. For example, in one embodiment, the appliance **200** comprises a first processor **262** and a second processor **262'**. In other embodiments, the processor **262** or **262'** comprises a multi-core processor.

Although the hardware layer **206** of appliance **200** is generally illustrated with an encryption processor **260**, processor **260** may be a processor for performing functions related to any encryption protocol, such as the Secure Socket Layer (SSL) or Transport Layer Security (TLS) protocol. In some embodiments, the processor **260** may be a general purpose processor (GPP), and in further embodiments, may have executable instructions for performing processing of any security related protocol.

Although the hardware layer **206** of appliance **200** is illustrated with certain elements in FIG. 2, the hardware portions or components of appliance **200** may comprise any type and form of elements, hardware or software, of a computing device, such as the computing device **100** illustrated and discussed herein in conjunction with FIGS. 1E and 1F. In some embodiments, the appliance **200** may comprise a server, gateway, router, switch, bridge or other type of computing or network device, and have any hardware and/or software elements associated therewith.

The operating system of appliance **200** allocates, manages, or otherwise segregates the available system memory into kernel space **204** and user space **204**. In example software architecture **200**, the operating system may be any type and/or form of Unix operating system although the invention is not so limited. As such, the appliance **200** can be running any operating system such as any of the versions of the Microsoft® Windows operating systems, the different releases of the Unix and Linux operating systems, any version of the Mac OS® for Macintosh computers, any embedded operating system, any network operating system, any real-time operating system, any open source operating system, any proprietary operating system, any operating systems for mobile computing devices or network devices, or any other operating system capable of running on the appliance **200** and performing the operations described herein.

The kernel space **204** is reserved for running the kernel **230**, including any device drivers, kernel extensions or other kernel related software. As known to those skilled in the art, the kernel **230** is the core of the operating system, and provides access, control, and management of resources and hardware-related elements of the application **104**. In accordance with an embodiment of the appliance **200**, the kernel space **204** also includes a number of network services or processes working in conjunction with a cache manager **232**, sometimes also referred to as the integrated cache, the benefits of which are described in detail further herein. Additionally, the embodiment of the kernel **230** will depend on the embodiment of the operating system installed, configured, or otherwise used by the device **200**.

In one embodiment, the device **200** comprises one network stack **267**, such as a TCP/IP based stack, for communicating with the client **102** and/or the server **106**. In one embodiment, the network stack **267** is used to communicate with a first network, such as network **108**, and a second network **110**. In some embodiments, the device **200** terminates a first transport layer connection, such as a TCP connection of a client **102**, and establishes a second transport layer connection to a server **106** for use by the client **102**, e.g., the second transport layer connection is terminated at the appliance **200** and the server **106**. The first and second transport layer connections may be established via a single network stack **267**. In other embodiments, the device **200** may comprise multiple network stacks, for example **267** and **267'**, and the first transport layer connection may be established or terminated at one network stack **267**, and the second transport layer connection on the second network stack **267'**. For example, one network stack may be for receiving and transmitting network packet on a first network, and another network stack for receiving and transmitting network packets on a second network. In one embodiment, the network stack **267** comprises a buffer **243** for queuing one or more network packets for transmission by the appliance **200**.

As shown in FIG. 2, the kernel space **204** includes the cache manager **232**, a high-speed layer 2-7 integrated packet engine **240**, an encryption engine **234**, a policy engine **236** and multi-protocol compression logic **238**. Running these components or processes **232**, **240**, **234**, **236** and **238** in kernel space **204** or kernel mode instead of the user space **202** improves the performance of each of these components, alone and in combination. Kernel operation means that these components or processes **232**, **240**, **234**, **236** and **238** run in the core address space of the operating system of the device **200**. For example, running the encryption engine **234** in kernel mode improves encryption performance by moving encryption and decryption operations to the kernel, thereby reducing the number of transitions between the memory space or a kernel thread in

kernel mode and the memory space or a thread in user mode. For example, data obtained in kernel mode may not need to be passed or copied to a process or thread running in user mode, such as from a kernel level data structure to a user level data structure. In another aspect, the number of context switches between kernel mode and user mode are also reduced. Additionally, synchronization of and communications between any of the components or processes **232**, **240**, **235**, **236** and **238** can be performed more efficiently in the kernel space **204**.

In some embodiments, any portion of the components **232**, **240**, **234**, **236** and **238** may run or operate in the kernel space **204**, while other portions of these components **232**, **240**, **234**, **236** and **238** may run or operate in user space **202**. In one embodiment, the appliance **200** uses a kernel-level data structure providing access to any portion of one or more network packets, for example, a network packet comprising a request from a client **102** or a response from a server **106**. In some embodiments, the kernel-level data structure may be obtained by the packet engine **240** via a transport layer driver interface or filter to the network stack **267**. The kernel-level data structure may comprise any interface and/or data accessible via the kernel space **204** related to the network stack **267**, network traffic or packets received or transmitted by the network stack **267**. In other embodiments, the kernel-level data structure may be used by any of the components or processes **232**, **240**, **234**, **236** and **238** to perform the desired operation of the component or process. In one embodiment, a component **232**, **240**, **234**, **236** and **238** is running in kernel mode **204** when using the kernel-level data structure, while in another embodiment, the component **232**, **240**, **234**, **236** and **238** is running in user mode when using the kernel-level data structure. In some embodiments, the kernel-level data structure may be copied or passed to a second kernel-level data structure, or any desired user-level data structure.

The cache manager **232** may comprise software, hardware or any combination of software and hardware to provide cache access, control and management of any type and form of content, such as objects or dynamically generated objects served by the originating servers **106**. The data, objects or content processed and stored by the cache manager **232** may comprise data in any format, such as a markup language, or communicated via any protocol. In some embodiments, the cache manager **232** duplicates original data stored elsewhere or data previously computed, generated or transmitted, in which the original data may require longer access time to fetch, compute or otherwise obtain relative to reading a cache memory element. Once the data is stored in the cache memory element, future use can be made by accessing the cached copy rather than refetching or recomputing the original data, thereby reducing the access time. In some embodiments, the cache memory element may comprise a data object in memory **264** of device **200**. In other embodiments, the cache memory element may comprise memory having a faster access time than memory **264**. In another embodiment, the cache memory element may comprise any type and form of storage element of the device **200**, such as a portion of a hard disk. In some embodiments, the processing unit **262** may provide cache memory for use by the cache manager **232**. In yet further embodiments, the cache manager **232** may use any portion and combination of memory, storage, or the processing unit for caching data, objects, and other content.

Furthermore, the cache manager **232** includes any logic, functions, rules, or operations to perform any embodiments of the techniques of the appliance **200** described herein. For example, the cache manager **232** includes logic or functionality to invalidate objects based on the expiration of an invalidity

23

duration time period or upon receipt of an invalidation command from a client **102** or server **106**. In some embodiments, the cache manager **232** may operate as a program, service, process or task executing in the kernel space **204**, and in other embodiments, in the user space **202**. In one embodiment, a first portion of the cache manager **232** executes in the user space **202** while a second portion executes in the kernel space **204**. In some embodiments, the cache manager **232** can comprise any type of general purpose processor (GPP), or any other type of integrated circuit, such as a Field Programmable Gate Array (FPGA), Programmable Logic Device (PLD), or Application Specific Integrated Circuit (ASIC).

The policy engine **236** may include, for example, an intelligent statistical engine or other programmable application (s). In one embodiment, the policy engine **236** provides a configuration mechanism to allow a user to identify, specify, define or configure a caching policy. Policy engine **236**, in some embodiments, also has access to memory to support data structures such as lookup tables or hash tables to enable user-selected caching policy decisions. In other embodiments, the policy engine **236** may comprise any logic, rules, functions or operations to determine and provide access, control and management of objects, data or content being cached by the appliance **200** in addition to access, control and management of security, network traffic, network access, compression or any other function or operation performed by the appliance **200**. Further examples of specific caching policies are further described herein.

The encryption engine **234** comprises any logic, business rules, functions or operations for handling the processing of any security related protocol, such as SSL or TLS, or any function related thereto. For example, the encryption engine **234** encrypts and decrypts network packets, or any portion thereof, communicated via the appliance **200**. The encryption engine **234** may also setup or establish SSL or TLS connections on behalf of the client **102a-102n**, server **106a-106n**, or appliance **200**. As such, the encryption engine **234** provides offloading and acceleration of SSL processing. In one embodiment, the encryption engine **234** uses a tunneling protocol to provide a virtual private network between a client **102a-102n** and a server **106a-106n**. In some embodiments, the encryption engine **234** is in communication with the Encryption processor **260**. In other embodiments, the encryption engine **234** comprises executable instructions running on the Encryption processor **260**.

The multi-protocol compression engine **238** comprises any logic, business rules, function or operations for compressing one or more protocols of a network packet, such as any of the protocols used by the network stack **267** of the device **200**. In one embodiment, multi-protocol compression engine **238** compresses bi-directionally between clients **102a-102n** and servers **106a-106n** any TCP/IP based protocol, including Messaging Application Programming Interface (MAPI) (email), File Transfer Protocol (FTP), HyperText Transfer Protocol (HTTP), Common Internet File System (CIFS) protocol (file transfer), Independent Computing Architecture (ICA) protocol, Remote Desktop Protocol (RDP), Wireless Application Protocol (WAP), Mobile IP protocol, and Voice Over IP (VoIP) protocol. In other embodiments, multi-protocol compression engine **238** provides compression of HyperText Markup Language (HTML) based protocols and in some embodiments, provides compression of any markup languages, such as the Extensible Markup Language (XML). In one embodiment, the multi-protocol compression engine **238** provides compression of any high-performance protocol, such as any protocol designed for appliance **200** to appliance **200** communications. In another embodiment, the multi-protocol

24

compression engine **238** compresses any payload of or any communication using a modified transport control protocol, such as Transaction TCP (T/TCP), TCP with selection acknowledgements (TCP-SACK), TCP with large windows (TCP-LW), a congestion prediction protocol such as the TCP-Vegas protocol, and a TCP spoofing protocol.

As such, the multi-protocol compression engine **238** accelerates performance for users accessing applications via desktop clients, e.g., Microsoft Outlook and non-Web thin clients, such as any client launched by popular enterprise applications like Oracle, SAP and Siebel, and even mobile clients, such as the Pocket PC. In some embodiments, the multi-protocol compression engine **238** by executing in the kernel mode **204** and integrating with packet processing engine **240** accessing the network stack **267** is able to compress any of the protocols carried by the TCP/IP protocol, such as any application layer protocol.

High speed layer **2-7** integrated packet engine **240**, also generally referred to as a packet processing engine or packet engine, is responsible for managing the kernel-level processing of packets received and transmitted by appliance **200** via network ports **266**. The high speed layer **2-7** integrated packet engine **240** may comprise a buffer for queuing one or more network packets during processing, such as for receipt of a network packet or transmission of a network packet. Additionally, the high speed layer **2-7** integrated packet engine **240** is in communication with one or more network stacks **267** to send and receive network packets via network ports **266**. The high speed layer **2-7** integrated packet engine **240** works in conjunction with encryption engine **234**, cache manager **232**, policy engine **236** and multi-protocol compression logic **238**. In particular, encryption engine **234** is configured to perform SSL processing of packets, policy engine **236** is configured to perform functions related to traffic management such as request-level content switching and request-level cache redirection, and multi-protocol compression logic **238** is configured to perform functions related to compression and decompression of data.

The high speed layer **2-7** integrated packet engine **240** includes a packet processing timer **242**. In one embodiment, the packet processing timer **242** provides one or more time intervals to trigger the processing of incoming, i.e., received, or outgoing, i.e., transmitted, network packets. In some embodiments, the high speed layer **2-7** integrated packet engine **240** processes network packets responsive to the timer **242**. The packet processing timer **242** provides any type and form of signal to the packet engine **240** to notify, trigger, or communicate a time related event, interval or occurrence. In many embodiments, the packet processing timer **242** operates in the order of milliseconds, such as for example 100 ms, 50 ms or 25 ms. For example, in some embodiments, the packet processing timer **242** provides time intervals or otherwise causes a network packet to be processed by the high speed layer **2-7** integrated packet engine **240** at a 10 ms time interval, while in other embodiments, at a 5 ms time interval, and still yet in further embodiments, as short as a 3, 2, or 1 ms time interval. The high speed layer **2-7** integrated packet engine **240** may be interfaced, integrated or in communication with the encryption engine **234**, cache manager **232**, policy engine **236** and multi-protocol compression engine **238** during operation. As such, any of the logic, functions, or operations of the encryption engine **234**, cache manager **232**, policy engine **236** and multi-protocol compression logic **238** may be performed responsive to the packet processing timer **242** and/or the packet engine **240**. Therefore, any of the logic, functions, or operations of the encryption engine **234**, cache manager **232**, policy engine **236** and multi-protocol compression

25

sion logic 238 may be performed at the granularity of time intervals provided via the packet processing timer 242, for example, at a time interval of less than or equal to 10 ms. For example, in one embodiment, the cache manager 232 may perform invalidation of any cached objects responsive to the high speed layer 2-7 integrated packet engine 240 and/or the packet processing timer 242. In another embodiment, the expiry or invalidation time of a cached object can be set to the same order of granularity as the time interval of the packet processing timer 242, such as at every 10 ms.

In contrast to kernel space 204, user space 202 is the memory area or portion of the operating system used by user mode applications or programs otherwise running in user mode. A user mode application may not access kernel space 204 directly and uses service calls in order to access kernel services. As shown in FIG. 2, user space 202 of appliance 200 includes a graphical user interface (GUI) 210, a command line interface (CLI) 212, shell services 214, health monitoring program 216, and daemon services 218. GUI 210 and CLI 212 provide a means by which a system administrator or other user can interact with and control the operation of appliance 200, such as via the operating system of the appliance 200. The GUI 210 or CLI 212 can comprise code running in user space 202 or kernel space 204. The GUI 210 may be any type and form of graphical user interface and may be presented via text, graphical or otherwise, by any type of program or application, such as a browser. The CLI 212 may be any type and form of command line or text-based interface, such as a command line provided by the operating system. For example, the CLI 212 may comprise a shell, which is a tool to enable users to interact with the operating system. In some embodiments, the CLI 212 may be provided via a bash, csh, tcsh, or ksh type shell. The shell services 214 comprises the programs, services, tasks, processes or executable instructions to support interaction with the appliance 200 or operating system by a user via the GUI 210 and/or CLI 212.

Health monitoring program 216 is used to monitor, check, report and ensure that network systems are functioning properly and that users are receiving requested content over a network. Health monitoring program 216 comprises one or more programs, services, tasks, processes or executable instructions to provide logic, rules, functions or operations for monitoring any activity of the appliance 200. In some embodiments, the health monitoring program 216 intercepts and inspects any network traffic passed via the appliance 200. In other embodiments, the health monitoring program 216 interfaces by any suitable means and/or mechanisms with one or more of the following: the encryption engine 234, cache manager 232, policy engine 236, multi-protocol compression logic 238, packet engine 240, daemon services 218, and shell services 214. As such, the health monitoring program 216 may call any application programming interface (API) to determine a state, status, or health of any portion of the appliance 200. For example, the health monitoring program 216 may ping or send a status inquiry on a periodic basis to check if a program, process, service or task is active and currently running. In another example, the health monitoring program 216 may check any status, error or history logs provided by any program, process, service or task to determine any condition, status or error with any portion of the appliance 200.

Daemon services 218 are programs that run continuously or in the background and handle periodic service requests received by appliance 200. In some embodiments, a daemon service may forward the requests to other programs or processes, such as another daemon service 218 as appropriate. As known to those skilled in the art, a daemon service 218 may

26

run unattended to perform continuous or periodic system wide functions, such as network control, or to perform any desired task. In some embodiments, one or more daemon services 218 run in the user space 202, while in other embodiments, one or more daemon services 218 run in the kernel space.

Referring now to FIG. 2B, another embodiment of the appliance 200 is depicted. In brief overview, the appliance 200 provides one or more of the following services, functionality or operations: SSL VPN connectivity 280, switching/load balancing 284, Domain Name Service resolution 286, acceleration 288 and an application firewall 290 for communications between one or more clients 102 and one or more servers 106. Each of the servers 106 may provide one or more network related services 270a-270n (referred to as services 270). For example, a server 106 may provide an http service 270. The appliance 200 comprises one or more virtual servers or virtual internet protocol servers, referred to as a vServer, VIP server, or just VIP 275a-275n (also referred herein as vServer 275). The vServer 275 receives, intercepts or otherwise processes communications between a client 102 and a server 106 in accordance with the configuration and operations of the appliance 200.

The vServer 275 may comprise software, hardware or any combination of software and hardware. The vServer 275 may comprise any type and form of program, service, task, process or executable instructions operating in user mode 202, kernel mode 204 or any combination thereof in the appliance 200. The vServer 275 includes any logic, functions, rules, or operations to perform any embodiments of the techniques described herein, such as SSL VPN 280, switching/load balancing 284, Domain Name Service resolution 286, acceleration 288 and an application firewall 290. In some embodiments, the vServer 275 establishes a connection to a service 270 of a server 106. The service 275 may comprise any program, application, process, task or set of executable instructions capable of connecting to and communicating to the appliance 200, client 102 or vServer 275. For example, the service 275 may comprise a web server, http server, ftp, email or database server. In some embodiments, the service 270 is a daemon process or network driver for listening, receiving and/or sending communications for an application, such as email, database or an enterprise application. In some embodiments, the service 270 may communicate on a specific IP address, or IP address and port.

In some embodiments, the vServer 275 applies one or more policies of the policy engine 236 to network communications between the client 102 and server 106. In one embodiment, the policies are associated with a vServer 275. In another embodiment, the policies are based on a user, or a group of users. In yet another embodiment, a policy is global and applies to one or more vServers 275a-275n, and any user or group of users communicating via the appliance 200. In some embodiments, the policies of the policy engine have conditions upon which the policy is applied based on any content of the communication, such as internet protocol address, port, protocol type, header or fields in a packet, or the context of the communication, such as user, group of the user, vServer 275, transport layer connection, and/or identification or attributes of the client 102 or server 106.

In other embodiments, the appliance 200 communicates or interfaces with the policy engine 236 to determine authentication and/or authorization of a remote user or a remote client 102 to access the computing environment 15, application, and/or data file from a server 106. In another embodiment, the appliance 200 communicates or interfaces with the policy engine 236 to determine authentication and/or authorization

27

of a remote user or a remote client **102** to have the application delivery system **190** deliver one or more of the computing environment **15**, application, and/or data file. In yet another embodiment, the appliance **200** establishes a VPN or SSL VPN connection based on the policy engine's **236** authentication and/or authorization of a remote user or a remote client **102**. In one embodiment, the appliance **200** controls the flow of network traffic and communication sessions based on policies of the policy engine **236**. For example, the appliance **200** may control the access to a computing environment **15**, application or data file based on the policy engine **236**.

In some embodiments, the vServer **275** establishes a transport layer connection, such as a TCP or UDP connection with a client **102** via the client agent **120**. In one embodiment, the vServer **275** listens for and receives communications from the client **102**. In other embodiments, the vServer **275** establishes a transport layer connection, such as a TCP or UDP connection with a client server **106**. In one embodiment, the vServer **275** establishes the transport layer connection to an internet protocol address and port of a server **270** running on the server **106**. In another embodiment, the vServer **275** associates a first transport layer connection to a client **102** with a second transport layer connection to the server **106**. In some embodiments, a vServer **275** establishes a pool of transport layer connections to a server **106** and multiplexes client requests via the pooled transport layer connections.

In some embodiments, the appliance **200** provides a SSL VPN connection **280** between a client **102** and a server **106**. For example, a client **102** on a first network **102** requests to establish a connection to a server **106** on a second network **104'**. In some embodiments, the second network **104'** is not routable from the first network **104**. In other embodiments, the client **102** is on a public network **104** and the server **106** is on a private network **104'**, such as a corporate network. In one embodiment, the client agent **120** intercepts communications of the client **102** on the first network **104**, encrypts the communications, and transmits the communications via a first transport layer connection to the appliance **200**. The appliance **200** associates the first transport layer connection on the first network **104** to a second transport layer connection to the server **106** on the second network **104**. The appliance **200** receives the intercepted communication from the client agent **102**, decrypts the communications, and transmits the communication to the server **106** on the second network **104** via the second transport layer connection. The second transport layer connection may be a pooled transport layer connection. As such, the appliance **200** provides an end-to-end secure transport layer connection for the client **102** between the two networks **104**, **104'**.

In one embodiment, the appliance **200** hosts an intranet internet protocol or IntranetIP **282** address of the client **102** on the virtual private network **104**. The client **102** has a local network identifier, such as an internet protocol (IP) address and/or host name on the first network **104**. When connected to the second network **104'** via the appliance **200**, the appliance **200** establishes, assigns or otherwise provides an IntranetIP address **282**, which is a network identifier, such as IP address and/or host name, for the client **102** on the second network **104'**. The appliance **200** listens for and receives on the second or private network **104'** for any communications directed towards the client **102** using the client's established IntranetIP **282**. In one embodiment, the appliance **200** acts as or on behalf of the client **102** on the second private network **104**. For example, in another embodiment, a vServer **275** listens for and responds to communications to the IntranetIP **282** of the client **102**. In some embodiments, if a computing device **100** on the second network **104'** transmits a request, the

28

appliance **200** processes the request as if it were the client **102**. For example, the appliance **200** may respond to a ping to the client's IntranetIP **282**. In another example, the appliance may establish a connection, such as a TCP or UDP connection, with computing device **100** on the second network **104** requesting a connection with the client's IntranetIP **282**.

In some embodiments, the appliance **200** provides one or more of the following acceleration techniques **288** to communications between the client **102** and server **106**: 1) compression; 2) decompression; 3) Transmission Control Protocol pooling; 4) Transmission Control Protocol multiplexing; 5) Transmission Control Protocol buffering; and 6) caching. In one embodiment, the appliance **200** relieves servers **106** of much of the processing load caused by repeatedly opening and closing transport layers connections to clients **102** by opening one or more transport layer connections with each server **106** and maintaining these connections to allow repeated data accesses by clients via the Internet. This technique is referred to herein as "connection pooling".

In some embodiments, in order to seamlessly splice communications from a client **102** to a server **106** via a pooled transport layer connection, the appliance **200** translates or multiplexes communications by modifying sequence number and acknowledgment numbers at the transport layer protocol level. This is referred to as "connection multiplexing". In some embodiments, no application layer protocol interaction is required. For example, in the case of an in-bound packet (that is, a packet received from a client **102**), the source network address of the packet is changed to that of an output port of appliance **200**, and the destination network address is changed to that of the intended server. In the case of an outbound packet (that is, one received from a server **106**), the source network address is changed from that of the server **106** to that of an output port of appliance **200** and the destination address is changed from that of appliance **200** to that of the requesting client **102**. The sequence numbers and acknowledgment numbers of the packet are also translated to sequence numbers and acknowledgement numbers expected by the client **102** on the appliance's **200** transport layer connection to the client **102**. In some embodiments, the packet checksum of the transport layer protocol is recalculated to account for these translations.

In another embodiment, the appliance **200** provides switching or load-balancing functionality **284** for communications between the client **102** and server **106**. In some embodiments, the appliance **200** distributes traffic and directs client requests to a server **106** based on layer **4** or application-layer request data. In one embodiment, although the network layer or layer **2** of the network packet identifies a destination server **106**, the appliance **200** determines the server **106** to distribute the network packet by application information and data carried as payload of the transport layer packet. In one embodiment, the health monitoring programs **216** of the appliance **200** monitor the health of servers to determine the server **106** for which to distribute a client's request. In some embodiments, if the appliance **200** detects a server **106** is not available or has a load over a predetermined threshold, the appliance **200** can direct or distribute client requests to another server **106**.

In some embodiments, the appliance **200** acts as a Domain Name Service (DNS) resolver or otherwise provides resolution of a DNS request from clients **102**. In some embodiments, the appliance intercepts a DNS request transmitted by the client **102**. In one embodiment, the appliance **200** responds to a client's DNS request with an IP address of or hosted by the appliance **200**. In this embodiment, the client **102** transmits network communication for the domain name

to the appliance 200. In another embodiment, the appliance 200 responds to a client's DNS request with an IP address of or hosted by a second appliance 200'. In some embodiments, the appliance 200 responds to a client's DNS request with an IP address of a server 106 determined by the appliance 200.

In yet another embodiment, the appliance 200 provides application firewall functionality 290 for communications between the client 102 and server 106. In one embodiment, the policy engine 236 provides rules for detecting and blocking illegitimate requests. In some embodiments, the application firewall 290 protects against denial of service (DoS) attacks. In other embodiments, the appliance inspects the content of intercepted requests to identify and block application-based attacks. In some embodiments, the rules/policy engine 236 comprises one or more application firewall or security control policies for providing protections against various classes and types of web or Internet based vulnerabilities, such as one or more of the following: 1) buffer overflow, 2) CGI-BIN parameter manipulation, 3) form/hidden field manipulation, 4) forceful browsing, 5) cookie or session poisoning, 6) broken access control list (ACLs) or weak passwords, 7) cross-site scripting (XSS), 8) command injection, 9) SQL injection, 10) error triggering sensitive information leak, 11) insecure use of cryptography, 12) server misconfiguration, 13) back doors and debug options, 14) website defacement, 15) platform or operating systems vulnerabilities, and 16) zero-day exploits. In an embodiment, the application firewall 290 provides HTML form field protection in the form of inspecting or analyzing the network communication for one or more of the following: 1) required fields are returned, 2) no added field allowed, 3) read-only and hidden field enforcement, 4) drop-down list and radio button field conformance, and 5) form-field max-length enforcement. In some embodiments, the application firewall 290 ensures cookies are not modified. In other embodiments, the application firewall 290 protects against forceful browsing by enforcing legal URLs.

In still yet other embodiments, the application firewall 290 protects any confidential information contained in the network communication. The application firewall 290 may inspect or analyze any network communication in accordance with the rules or policies of the engine 236 to identify any confidential information in any field of the network packet. In some embodiments, the application firewall 290 identifies in the network communication one or more occurrences of a credit card number, password, social security number, name, patient code, contact information, and age. The encoded portion of the network communication may comprise these occurrences or the confidential information. Based on these occurrences, in one embodiment, the application firewall 290 may take a policy action on the network communication, such as prevent transmission of the network communication. In another embodiment, the application firewall 290 may rewrite, remove or otherwise mask such identified occurrence or confidential information.

Still referring to FIG. 2B, the appliance 200 may include a performance monitoring agent 197 as discussed above in conjunction with FIG. 1D. In one embodiment, the appliance 200 receives the monitoring agent 197 from the monitoring service 198 or monitoring server 106 as depicted in FIG. 1D. In some embodiments, the appliance 200 stores the monitoring agent 197 in storage, such as disk, for delivery to any client or server in communication with the appliance 200. For example, in one embodiment, the appliance 200 transmits the monitoring agent 197 to a client upon receiving a request to establish a transport layer connection. In other embodiments, the appliance 200 transmits the monitoring agent 197 upon

establishing the transport layer connection with the client 102. In another embodiment, the appliance 200 transmits the monitoring agent 197 to the client upon intercepting or detecting a request for a web page. In yet another embodiment, the appliance 200 transmits the monitoring agent 197 to a client or a server in response to a request from the monitoring server 198. In one embodiment, the appliance 200 transmits the monitoring agent 197 to a second appliance 200' or appliance 205.

In other embodiments, the appliance 200 executes the monitoring agent 197. In one embodiment, the monitoring agent 197 measures and monitors the performance of any application, program, process, service, task or thread executing on the appliance 200. For example, the monitoring agent 197 may monitor and measure performance and operation of vServers 275A-275N. In another embodiment, the monitoring agent 197 measures and monitors the performance of any transport layer connections of the appliance 200. In some embodiments, the monitoring agent 197 measures and monitors the performance of any user sessions traversing the appliance 200. In one embodiment, the monitoring agent 197 measures and monitors the performance of any virtual private network connections and/or sessions traversing the appliance 200, such as an SSL VPN session. In still further embodiments, the monitoring agent 197 measures and monitors the memory, CPU and disk usage and performance of the appliance 200. In yet another embodiment, the monitoring agent 197 measures and monitors the performance of any acceleration technique 288 performed by the appliance 200, such as SSL offloading, connection pooling and multiplexing, caching, and compression. In some embodiments, the monitoring agent 197 measures and monitors the performance of any load balancing and/or content switching 284 performed by the appliance 200. In other embodiments, the monitoring agent 197 measures and monitors the performance of application firewall 290 protection and processing performed by the appliance 200.

C. Client Agent

Referring now to FIG. 3, an embodiment of the client agent 120 is depicted. The client 102 includes a client agent 120 for establishing and exchanging communications with the appliance 200 and/or server 106 via a network 104. In brief overview, the client 102 operates on computing device 100 having an operating system with a kernel mode 302 and a user mode 303, and a network stack 310 with one or more layers 310a-310b. The client 102 may have installed and/or execute one or more applications. In some embodiments, one or more applications may communicate via the network stack 310 to a network 104. One of the applications, such as a web browser, may also include a first program 322. For example, the first program 322 may be used in some embodiments to install and/or execute the client agent 120, or any portion thereof. The client agent 120 includes an interception mechanism, or interceptor 350, for intercepting network communications from the network stack 310 from the one or more applications.

The network stack 310 of the client 102 may comprise any type and form of software, or hardware, or any combinations thereof, for providing connectivity to and communications with a network. In one embodiment, the network stack 310 comprises a software implementation for a network protocol suite. The network stack 310 may comprise one or more network layers, such as any networks layers of the Open Systems Interconnection (OSI) communications model as those skilled in the art recognize and appreciate. As such, the network stack 310 may comprise any type and form of protocols for any of the following layers of the OSI model: 1)

31

physical link layer, 2) data link layer, 3) network layer, 4) transport layer, 5) session layer, 6) presentation layer, and 7) application layer. In one embodiment, the network stack **310** may comprise a transport control protocol (TCP) over the network layer protocol of the internet protocol (IP), generally referred to as TCP/IP. In some embodiments, the TCP/IP protocol may be carried over the Ethernet protocol, which may comprise any of the family of IEEE wide-area-network (WAN) or local-area-network (LAN) protocols, such as those protocols covered by the IEEE 802.3. In some embodiments, the network stack **310** comprises any type and form of a wireless protocol, such as IEEE 802.11 and/or mobile internet protocol.

In view of a TCP/IP based network, any TCP/IP based protocol may be used, including Messaging Application Programming Interface (MAPI) (email), File Transfer Protocol (FTP), HyperText Transfer Protocol (HTTP), Common Internet File System (CIFS) protocol (file transfer), Independent Computing Architecture (ICA) protocol, Remote Desktop Protocol (RDP), Wireless Application Protocol (WAP), Mobile IP protocol, and Voice Over IP (VoIP) protocol. In another embodiment, the network stack **310** comprises any type and form of transport control protocol, such as a modified transport control protocol, for example a Transaction TCP (T/TCP), TCP with selection acknowledgements (TCP-SACK), TCP with large windows (TCP-LW), a congestion prediction protocol such as the TCP-Vegas protocol, and a TCP spoofing protocol. In other embodiments, any type and form of user datagram protocol (UDP), such as UDP over IP, may be used by the network stack **310**, such as for voice communications or real-time data communications.

Furthermore, the network stack **310** may include one or more network drivers supporting the one or more layers, such as a TCP driver or a network layer driver. The network drivers may be included as part of the operating system of the computing device **100** or as part of any network interface cards or other network access components of the computing device **100**. In some embodiments, any of the network drivers of the network stack **310** may be customized, modified or adapted to provide a custom or modified portion of the network stack **310** in support of any of the techniques described herein. In other embodiments, the acceleration program **302** is designed and constructed to operate with or work in conjunction with the network stack **310** installed or otherwise provided by the operating system of the client **102**.

The network stack **310** comprises any type and form of interfaces for receiving, obtaining, providing or otherwise accessing any information and data related to network communications of the client **102**. In one embodiment, an interface to the network stack **310** comprises an application programming interface (API). The interface may also comprise any function call, hooking or filtering mechanism, event or call back mechanism, or any type of interfacing technique. The network stack **310** via the interface may receive or provide any type and form of data structure, such as an object, related to functionality or operation of the network stack **310**. For example, the data structure may comprise information and data related to a network packet or one or more network packets. In some embodiments, the data structure comprises a portion of the network packet processed at a protocol layer of the network stack **310**, such as a network packet of the transport layer. In some embodiments, the data structure **325** comprises a kernel-level data structure, while in other embodiments, the data structure **325** comprises a user-mode data structure. A kernel-level data structure may comprise a data structure obtained or related to a portion of the network stack **310** operating in kernel-mode **302**, or a network driver

32

or other software running in kernel-mode **302**, or any data structure obtained or received by a service, process, task, thread or other executable instructions running or operating in kernel-mode of the operating system.

Additionally, some portions of the network stack **310** may execute or operate in kernel-mode **302**, for example, the data link or network layer, while other portions execute or operate in user-mode **303**, such as an application layer of the network stack **310**. For example, a first portion **310a** of the network stack may provide user-mode access to the network stack **310** to an application while a second portion **310b** of the network stack **310** provides access to a network. In some embodiments, a first portion **310a** of the network stack may comprise one or more upper layers of the network stack **310**, such as any of layers **5-7**. In other embodiments, a second portion **310b** of the network stack **310** comprises one or more lower layers, such as any of layers **1-4**. Each of the first portion **310a** and second portion **310b** of the network stack **310** may comprise any portion of the network stack **310**, at any one or more network layers, in user-mode **203**, kernel-mode, **202**, or combinations thereof, or at any portion of a network layer or interface point to a network layer or any portion of or interface point to the user-mode **203** and kernel-mode **203**.

The interceptor **350** may comprise software, hardware, or any combination of software and hardware. In one embodiment, the interceptor **350** intercept a network communication at any point in the network stack **310**, and redirects or transmits the network communication to a destination desired, managed or controlled by the interceptor **350** or client agent **120**. For example, the interceptor **350** may intercept a network communication of a network stack **310** of a first network and transmit the network communication to the appliance **200** for transmission on a second network **104**. In some embodiments, the interceptor **350** comprises any type of interceptor **350** comprises a driver, such as a network driver constructed and designed to interface and work with the network stack **310**. In some embodiments, the client agent **120** and/or interceptor **350** operates at one or more layers of the network stack **310**, such as at the transport layer. In one embodiment, the interceptor **350** comprises a filter driver, hooking mechanism, or any form and type of suitable network driver interface that interfaces to the transport layer of the network stack, such as via the transport driver interface (TDI). In some embodiments, the interceptor **350** interfaces to a first protocol layer, such as the transport layer and another protocol layer, such as any layer above the transport protocol layer, for example, an application protocol layer. In one embodiment, the interceptor **350** may comprise a driver complying with the Network Driver Interface Specification (NDIS), or a NDIS driver. In another embodiment, the interceptor **350** may comprise a mini-filter or a mini-port driver. In one embodiment, the interceptor **350**, or portion thereof, operates in kernel-mode **202**. In another embodiment, the interceptor **350**, or portion thereof, operates in user-mode **203**. In some embodiments, a portion of the interceptor **350** operates in kernel-mode **202** while another portion of the interceptor **350** operates in user-mode **203**. In other embodiments, the client agent **120** operates in user-mode **203** but interfaces via the interceptor **350** to a kernel-mode driver, process, service, task or portion of the operating system, such as to obtain a kernel-level data structure **225**. In further embodiments, the interceptor **350** is a user-mode application or program, such as application.

In one embodiment, the interceptor **350** intercepts any transport layer connection requests. In these embodiments, the interceptor **350** execute transport layer application programming interface (API) calls to set the destination information, such as destination IP address and/or port to a desired

33

location for the location. In this manner, the interceptor **350** intercepts and redirects the transport layer connection to a IP address and port controlled or managed by the interceptor **350** or client agent **120**. In one embodiment, the interceptor **350** sets the destination information for the connection to a local IP address and port of the client **102** on which the client agent **120** is listening. For example, the client agent **120** may comprise a proxy service listening on a local IP address and port for redirected transport layer communications. In some embodiments, the client agent **120** then communicates the redirected transport layer communication to the appliance **200**.

In some embodiments, the interceptor **350** intercepts a Domain Name Service (DNS) request. In one embodiment, the client agent **120** and/or interceptor **350** resolves the DNS request. In another embodiment, the interceptor transmits the intercepted DNS request to the appliance **200** for DNS resolution. In one embodiment, the appliance **200** resolves the DNS request and communicates the DNS response to the client agent **120**. In some embodiments, the appliance **200** resolves the DNS request via another appliance **200'** or a DNS server **106**.

In yet another embodiment, the client agent **120** may comprise two agents **120** and **120'**. In one embodiment, a first agent **120** may comprise an interceptor **350** operating at the network layer of the network stack **310**. In some embodiments, the first agent **120** intercepts network layer requests such as Internet Control Message Protocol (ICMP) requests (e.g., ping and traceroute). In other embodiments, the second agent **120'** may operate at the transport layer and intercept transport layer communications. In some embodiments, the first agent **120** intercepts communications at one layer of the network stack **210** and interfaces with or communicates the intercepted communication to the second agent **120'**.

The client agent **120** and/or interceptor **350** may operate at or interface with a protocol layer in a manner transparent to any other protocol layer of the network stack **310**. For example, in one embodiment, the interceptor **350** operates or interfaces with the transport layer of the network stack **310** transparently to any protocol layer below the transport layer, such as the network layer, and any protocol layer above the transport layer, such as the session, presentation or application layer protocols. This allows the other protocol layers of the network stack **310** to operate as desired and without modification for using the interceptor **350**. As such, the client agent **120** and/or interceptor **350** can interface with the transport layer to secure, optimize, accelerate, route or load-balance any communications provided via any protocol carried by the transport layer, such as any application layer protocol over TCP/IP.

Furthermore, the client agent **120** and/or interceptor may operate at or interface with the network stack **310** in a manner transparent to any application, a user of the client **102**, and any other computing device, such as a server, in communications with the client **102**. The client agent **120** and/or interceptor **350** may be installed and/or executed on the client **102** in a manner without modification of an application. In some embodiments, the user of the client **102** or a computing device in communications with the client **102** are not aware of the existence, execution or operation of the client agent **120** and/or interceptor **350**. As such, in some embodiments, the client agent **120** and/or interceptor **350** is installed, executed, and/or operated transparently to an application, user of the client **102**, another computing device, such as a server, or any of the protocol layers above and/or below the protocol layer interfaced to by the interceptor **350**.

34

The client agent **120** includes an acceleration program **302**, a streaming client **306**, a collection agent **304**, and/or monitoring agent **197**. In one embodiment, the client agent **120** comprises an Independent Computing Architecture (ICA) client, or any portion thereof, developed by Citrix Systems, Inc. of Fort Lauderdale, Fla., and is also referred to as an ICA client. In some embodiments, the client **120** comprises an application streaming client **306** for streaming an application from a server **106** to a client **102**. In some embodiments, the client agent **120** comprises an acceleration program **302** for accelerating communications between client **102** and server **106**. In another embodiment, the client agent **120** includes a collection agent **304** for performing end-point detection/ scanning and collecting end-point information for the appliance **200** and/or server **106**.

In some embodiments, the acceleration program **302** comprises a client-side acceleration program for performing one or more acceleration techniques to accelerate, enhance or otherwise improve a client's communications with and/or access to a server **106**, such as accessing an application provided by a server **106**. The logic, functions, and/or operations of the executable instructions of the acceleration program **302** may perform one or more of the following acceleration techniques: 1) multi-protocol compression, 2) transport control protocol pooling, 3) transport control protocol multiplexing, 4) transport control protocol buffering, and 5) caching via a cache manager. Additionally, the acceleration program **302** may perform encryption and/or decryption of any communications received and/or transmitted by the client **102**. In some embodiments, the acceleration program **302** performs one or more of the acceleration techniques in an integrated manner or fashion. Additionally, the acceleration program **302** can perform compression on any of the protocols, or multiple protocols, carried as a payload of a network packet of the transport layer protocol.

The streaming client **306** comprises an application, program, process, service, task or executable instructions for receiving and executing a streamed application from a server **106**. A server **106** may stream one or more application data files to the streaming client **306** for playing, executing or otherwise causing to be executed the application on the client **102**. In some embodiments, the server **106** transmits a set of compressed or packaged application data files to the streaming client **306**. In some embodiments, the plurality of application files are compressed and stored on a file server within an archive file such as a CAB, ZIP, SIT, TAR, JAR or other archive. In one embodiment, the server **106** decompresses, unpackages or unarchives the application files and transmits the files to the client **102**. In another embodiment, the client **102** decompresses, unpackages or unarchives the application files. The streaming client **306** dynamically installs the application, or portion thereof, and executes the application. In one embodiment, the streaming client **306** may be an executable program. In some embodiments, the streaming client **306** may be able to launch another executable program.

The collection agent **304** comprises an application, program, process, service, task or executable instructions for identifying, obtaining and/or collecting information about the client **102**. In some embodiments, the appliance **200** transmits the collection agent **304** to the client **102** or client agent **120**. The collection agent **304** may be configured according to one or more policies of the policy engine **236** of the appliance. In other embodiments, the collection agent **304** transmits collected information on the client **102** to the appliance **200**. In one embodiment, the policy engine **236** of the appliance **200** uses the collected information to determine and provide

access, authentication and authorization control of the client's connection to a network **104**.

In one embodiment, the collection agent **304** comprises an end-point detection and scanning mechanism, which identifies and determines one or more attributes or characteristics of the client. For example, the collection agent **304** may identify and determine any one or more of the following client-side attributes: 1) the operating system and/or a version of an operating system, 2) a service pack of the operating system, 3) a running service, 4) a running process, and 5) a file. The collection agent **304** may also identify and determine the presence or versions of any one or more of the following on the client: 1) antivirus software, 2) personal firewall software, 3) anti-spam software, and 4) internet security software. The policy engine **236** may have one or more policies based on any one or more of the attributes or characteristics of the client or client-side attributes.

In some embodiments, the client agent **120** includes a monitoring agent **197** as discussed in conjunction with FIGS. 1D and 2B. The monitoring agent **197** may be any type and form of script, such as Visual Basic or Java script. In one embodiment, the monitoring agent **197** monitors and measures performance of any portion of the client agent **120**. For example, in some embodiments, the monitoring agent **197** monitors and measures performance of the acceleration program **302**. In another embodiment, the monitoring agent **197** monitors and measures performance of the streaming client **306**. In other embodiments, the monitoring agent **197** monitors and measures performance of the collection agent **304**. In still another embodiment, the monitoring agent **197** monitors and measures performance of the interceptor **350**. In some embodiments, the monitoring agent **197** monitors and measures any resource of the client **102**, such as memory, CPU and disk.

The monitoring agent **197** may monitor and measure performance of any application of the client. In one embodiment, the monitoring agent **197** monitors and measures performance of a browser on the client **102**. In some embodiments, the monitoring agent **197** monitors and measures performance of any application delivered via the client agent **120**. In other embodiments, the monitoring agent **197** measures and monitors end user response times for an application, such as web-based or HTTP response times. The monitoring agent **197** may monitor and measure performance of an ICA or RDP client. In another embodiment, the monitoring agent **197** measures and monitors metrics for a user session or application session. In some embodiments, monitoring agent **197** measures and monitors an ICA or RDP session. In one embodiment, the monitoring agent **197** measures and monitors the performance of the appliance **200** in accelerating delivery of an application and/or data to the client **102**.

In some embodiments and still referring to FIG. 3, a first program **322** may be used to install and/or execute the client agent **120**, or portion thereof, such as the interceptor **350**, automatically, silently, transparently, or otherwise. In one embodiment, the first program **322** comprises a plugin component, such as an ActiveX control or Java control or script that is loaded into and executed by an application. For example, the first program comprises an ActiveX control loaded and run by a web browser application, such as in the memory space or context of the application. In another embodiment, the first program **322** comprises a set of executable instructions loaded into and run by the application, such as a browser. In one embodiment, the first program **322** comprises a designed and constructed program to install the client agent **120**. In some embodiments, the first program **322** obtains, downloads, or receives the client agent **120** via the network

from another computing device. In another embodiment, the first program **322** is an installer program or a plug and play manager for installing programs, such as network drivers, on the operating system of the client **102**.

D. Systems and Methods for Providing Virtualized Application Delivery Controller

Referring now to FIG. 4A, a block diagram depicts one embodiment of a virtualization environment **400**. In brief overview, a computing device **100** includes a hypervisor layer, a virtualization layer, and a hardware layer. The hypervisor layer includes a hypervisor **401** (also referred to as a virtualization manager) that allocates and manages access to a number of physical resources in the hardware layer (e.g., the processor(s) **421**, and disk(s) **428**) by at least one virtual machine executing in the virtualization layer. The virtualization layer includes at least one operating system **410** and a plurality of virtual resources allocated to the at least one operating system **410**. Virtual resources may include, without limitation, a plurality of virtual processors **432a**, **432b**, **432c** (generally **432**), and virtual disks **442a**, **442b**, **442c** (generally **442**), as well as virtual resources such as virtual memory and virtual network interfaces. The plurality of virtual resources and the operating system **410** may be referred to as a virtual machine **406**. A virtual machine **406** may include a control operating system **405** in communication with the hypervisor **401** and used to execute applications for managing and configuring other virtual machines on the computing device **100**.

In greater detail, a hypervisor **401** may provide virtual resources to an operating system in any manner which simulates the operating system having access to a physical device. A hypervisor **401** may provide virtual resources to any number of guest operating systems **410a**, **410b** (generally **410**). In some embodiments, a computing device **100** executes one or more types of hypervisors. In these embodiments, hypervisors may be used to emulate virtual hardware, partition physical hardware, virtualize physical hardware, and execute virtual machines that provide access to computing environments. Hypervisors may include those manufactured by VMWare, Inc., of Palo Alto, Calif.; the XEN hypervisor, an open source product whose development is overseen by the open source Xen.org community; HyperV, VirtualServer or virtual PC hypervisors provided by Microsoft, or others. In some embodiments, a computing device **100** executing a hypervisor that creates a virtual machine platform on which guest operating systems may execute is referred to as a host server. In one of these embodiments, for example, the computing device **100** is a XEN SERVER provided by Citrix Systems, Inc., of Fort Lauderdale, Fla.

In some embodiments, a hypervisor **401** executes within an operating system executing on a computing device. In one of these embodiments, a computing device executing an operating system and a hypervisor **401** may be said to have a host operating system (the operating system executing on the computing device), and a guest operating system (an operating system executing within a computing resource partition provided by the hypervisor **401**). In other embodiments, a hypervisor **401** interacts directly with hardware on a computing device, instead of executing on a host operating system. In one of these embodiments, the hypervisor **401** may be said to be executing on "bare metal," referring to the hardware comprising the computing device.

In some embodiments, a hypervisor **401** may create a virtual machine **406a-c** (generally **406**) in which an operating system **410** executes. In one of these embodiments, for example, the hypervisor **401** loads a virtual machine image to create a virtual machine **406**. In another of these embodiments, the hypervisor **401** executes an operating system **410**

within the virtual machine **406**. In still another of these embodiments, the virtual machine **406** executes an operating system **410**.

In some embodiments, the hypervisor **401** controls processor scheduling and memory partitioning for a virtual machine **406** executing on the computing device **100**. In one of these embodiments, the hypervisor **401** controls the execution of at least one virtual machine **406**. In another of these embodiments, the hypervisor **401** presents at least one virtual machine **406** with an abstraction of at least one hardware resource provided by the computing device **100**. In other embodiments, the hypervisor **401** controls whether and how physical processor capabilities are presented to the virtual machine **406**.

A control operating system **405** may execute at least one application for managing and configuring the guest operating systems. In one embodiment, the control operating system **405** may execute an administrative application, such as an application including a user interface providing administrators with access to functionality for managing the execution of a virtual machine, including functionality for executing a virtual machine, terminating an execution of a virtual machine, or identifying a type of physical resource for allocation to the virtual machine. In another embodiment, the hypervisor **401** executes the control operating system **405** within a virtual machine **406** created by the hypervisor **401**. In still another embodiment, the control operating system **405** executes in a virtual machine **406** that is authorized to directly access physical resources on the computing device **100**. In some embodiments, a control operating system **405a** on a computing device **100a** may exchange data with a control operating system **405b** on a computing device **100b**, via communications between a hypervisor **401a** and a hypervisor **401b**. In this way, one or more computing devices **100** may exchange data with one or more of the other computing devices **100** regarding processors and other physical resources available in a pool of resources. In one of these embodiments, this functionality allows a hypervisor to manage a pool of resources distributed across a plurality of physical computing devices. In another of these embodiments, multiple hypervisors manage one or more of the guest operating systems executed on one of the computing devices **100**.

In one embodiment, the control operating system **405** executes in a virtual machine **406** that is authorized to interact with at least one guest operating system **410**. In another embodiment, a guest operating system **410** communicates with the control operating system **405** via the hypervisor **401** in order to request access to a disk or a network. In still another embodiment, the guest operating system **410** and the control operating system **405** may communicate via a communication channel established by the hypervisor **401**, such as, for example, via a plurality of shared memory pages made available by the hypervisor **401**.

In some embodiments, the control operating system **405** includes a network back-end driver for communicating directly with networking hardware provided by the computing device **100**. In one of these embodiments, the network back-end driver processes at least one virtual machine request from at least one guest operating system **110**. In other embodiments, the control operating system **405** includes a block back-end driver for communicating with a storage element on the computing device **100**. In one of these embodiments, the block back-end driver reads and writes data from the storage element based upon at least one request received from a guest operating system **410**.

In one embodiment, the control operating system **405** includes a tools stack **404**. In another embodiment, a tools

stack **404** provides functionality for interacting with the hypervisor **401**, communicating with other control operating systems **405** (for example, on a second computing device **100b**), or managing virtual machines **406b**, **406c** on the computing device **100**. In another embodiment, the tools stack **404** includes customized applications for providing improved management functionality to an administrator of a virtual machine farm. In some embodiments, at least one of the tools stack **404** and the control operating system **405** include a management API that provides an interface for remotely configuring and controlling virtual machines **406** running on a computing device **100**. In other embodiments, the control operating system **405** communicates with the hypervisor **401** through the tools stack **404**.

In one embodiment, the hypervisor **401** executes a guest operating system **410** within a virtual machine **406** created by the hypervisor **401**. In another embodiment, the guest operating system **410** provides a user of the computing device **100** with access to resources within a computing environment. In still another embodiment, a resource includes a program, an application, a document, a file, a plurality of applications, a plurality of files, an executable program file, a desktop environment, a computing environment, or other resource made available to a user of the computing device **100**. In yet another embodiment, the resource may be delivered to the computing device **100** via a plurality of access methods including, but not limited to, conventional installation directly on the computing device **100**, delivery to the computing device **100** via a method for application streaming, delivery to the computing device **100** of output data generated by an execution of the resource on a second computing device **100'** and communicated to the computing device **100** via a presentation layer protocol, delivery to the computing device **100** of output data generated by an execution of the resource via a virtual machine executing on a second computing device **100'**, or execution from a removable storage device connected to the computing device **100**, such as a USB device, or via a virtual machine executing on the computing device **100** and generating output data. In some embodiments, the computing device **100** transmits output data generated by the execution of the resource to another computing device **100'**.

In one embodiment, the guest operating system **410**, in conjunction with the virtual machine on which it executes, forms a fully-virtualized virtual machine which is not aware that it is a virtual machine; such a machine may be referred to as a "Domain U HVM (Hardware Virtual Machine) virtual machine". In another embodiment, a fully-virtualized machine includes software emulating a Basic Input/Output System (BIOS) in order to execute an operating system within the fully-virtualized machine. In still another embodiment, a fully-virtualized machine may include a driver that provides functionality by communicating with the hypervisor **401**. In such an embodiment, the driver may be aware that it executes within a virtualized environment. In another embodiment, the guest operating system **410**, in conjunction with the virtual machine on which it executes, forms a paravirtualized virtual machine, which is aware that it is a virtual machine; such a machine may be referred to as a "Domain U PV virtual machine". In another embodiment, a paravirtualized machine includes additional drivers that a fully-virtualized machine does not include. In still another embodiment, the paravirtualized machine includes the network back-end driver and the block back-end driver included in a control operating system **405**, as described above.

Referring now to FIG. 4B, a block diagram depicts one embodiment of a plurality of networked computing devices in a system in which at least one physical host executes a virtual

machine. In brief overview, the system includes a management component **404** and a hypervisor **401**. The system includes a plurality of computing devices **100**, a plurality of virtual machines **406**, a plurality of hypervisors **401**, a plurality of management components referred to variously as tools stacks **404** or management components **404**, and a physical resource **421**, **428**. The plurality of physical machines **100** may each be provided as computing devices **100**, described above in connection with FIGS. 1E-1H and 4A.

In greater detail, a physical disk **428** is provided by a computing device **100** and stores at least a portion of a virtual disk **442**. In some embodiments, a virtual disk **442** is associated with a plurality of physical disks **428**. In one of these embodiments, one or more computing devices **100** may exchange data with one or more of the other computing devices **100** regarding processors and other physical resources available in a pool of resources, allowing a hypervisor to manage a pool of resources distributed across a plurality of physical computing devices. In some embodiments, a computing device **100** on which a virtual machine **406** executes is referred to as a physical host **100** or as a host machine **100**.

The hypervisor executes on a processor on the computing device **100**. The hypervisor allocates, to a virtual disk, an amount of access to the physical disk. In one embodiment, the hypervisor **401** allocates an amount of space on the physical disk. In another embodiment, the hypervisor **401** allocates a plurality of pages on the physical disk. In some embodiments, the hypervisor provisions the virtual disk **442** as part of a process of initializing and executing a virtual machine **450**.

In one embodiment, the management component **404a** is referred to as a pool management component **404a**. In another embodiment, a management operating system **405a**, which may be referred to as a control operating system **405a**, includes the management component. In some embodiments, the management component is referred to as a tools stack. In one of these embodiments, the management component is the tools stack **404** described above in connection with FIG. 4A. In other embodiments, the management component **404** provides a user interface for receiving, from a user such as an administrator, an identification of a virtual machine **406** to provision and/or execute. In still other embodiments, the management component **404** provides a user interface for receiving, from a user such as an administrator, the request for migration of a virtual machine **406b** from one physical machine **100** to another. In further embodiments, the management component **404a** identifies a computing device **100b** on which to execute a requested virtual machine **406d** and instructs the hypervisor **401b** on the identified computing device **100b** to execute the identified virtual machine; such a management component may be referred to as a pool management component.

Referring now to FIG. 4C, embodiments of a virtual application delivery controller or virtual appliance **450** are depicted. In brief overview, any of the functionality and/or embodiments of the appliance **200** (e.g., an application delivery controller) described above in connection with FIGS. 2A and 2B may be deployed in any embodiment of the virtualized environment described above in connection with FIGS. 4A and 4B. Instead of the functionality of the application delivery controller being deployed in the form of an appliance **200**, such functionality may be deployed in a virtualized environment **400** on any computing device **100**, such as a client **102**, server **106** or appliance **200**.

Referring now to FIG. 4C, a diagram of an embodiment of a virtual appliance **450** operating on a hypervisor **401** of a

server **106** is depicted. As with the appliance **200** of FIGS. 2A and 2B, the virtual appliance **450** may provide functionality for availability, performance, offload and security. For availability, the virtual appliance may perform load balancing between layers **4** and **7** of the network and may also perform intelligent service health monitoring. For performance increases via network traffic acceleration, the virtual appliance may perform caching and compression. To offload processing of any servers, the virtual appliance may perform connection multiplexing and pooling and/or SSL processing. For security, the virtual appliance may perform any of the application firewall functionality and SSL VPN function of appliance **200**.

Any of the modules of the appliance **200** as described in connection with FIG. 2A may be packaged, combined, designed or constructed in a form of the virtualized appliance delivery controller **450** deployable as one or more software modules or components executable in a virtualized environment **300** or non-virtualized environment on any server, such as an off the shelf server. For example, the virtual appliance may be provided in the form of an installation package to install on a computing device. With reference to FIG. 2A, any of the cache manager **232**, policy engine **236**, compression **238**, encryption engine **234**, packet engine **240**, GUI **210**, CLI **212**, shell services **214** and health monitoring programs **216** may be designed and constructed as a software component or module to run on any operating system of a computing device and/or of a virtualized environment **300**. Instead of using the encryption processor **260**, processor **262**, memory **264** and network stack **267** of the appliance **200**, the virtualized appliance **400** may use any of these resources as provided by the virtualized environment **400** or as otherwise available on the server **106**.

Still referring to FIG. 4C, and in brief overview, any one or more vServers **275A-275N** may be in operation or executed in a virtualized environment **400** of any type of computing device **100**, such as any server **106**. Any of the modules or functionality of the appliance **200** described in connection with FIG. 2B may be designed and constructed to operate in either a virtualized or non-virtualized environment of a server. Any of the vServer **275**, SSL VPN **280**, Intranet UP **282**, Switching **284**, DNS **286**, acceleration **288**, App FW **280** and monitoring agent may be packaged, combined, designed or constructed in a form of application delivery controller **450** deployable as one or more software modules or components executable on a device and/or virtualized environment **400**.

In some embodiments, a server may execute multiple virtual machines **406a-406n** in the virtualization environment with each virtual machine running the same or different embodiments of the virtual application delivery controller **450**. In some embodiments, the server may execute one or more virtual appliances **450** on one or more virtual machines on a core of a multi-core processing system. In some embodiments, the server may execute one or more virtual appliances **450** on one or more virtual machines on each processor of a multiple processor device.

E. Systems and Methods for Providing A Multi-Core Architecture

In accordance with Moore's Law, the number of transistors that may be placed on an integrated circuit may double approximately every two years. However, CPU speed increases may reach plateaus, for example CPU speed has been around 3.5-4 GHz range since 2005. In some cases, CPU manufacturers may not rely on CPU speed increases to gain additional performance. Some CPU manufacturers may add additional cores to their processors to provide additional performance. Products, such as those of software and network-

41

ing vendors, that rely on CPUs for performance gains may improve their performance by leveraging these multi-core CPUs. The software designed and constructed for a single CPU may be redesigned and/or rewritten to take advantage of a multi-threaded, parallel architecture or otherwise a multi-core architecture.

A multi-core architecture of the appliance **200**, referred to as nCore or multi-core technology, allows the appliance in some embodiments to break the single core performance barrier and to leverage the power of multi-core CPUs. In the previous architecture described in connection with FIG. 2A, a single network or packet engine is run. The multiple cores of the nCore technology and architecture allow multiple packet engines to run concurrently and/or in parallel. With a packet engine running on each core, the appliance architecture leverages the processing capacity of additional cores. In some embodiments, this provides up to a 7× increase in performance and scalability.

Illustrated in FIG. 5A are some embodiments of work, task, load or network traffic distribution across one or more processor cores according to a type of parallelism or parallel computing scheme, such as functional parallelism, data parallelism or flow-based data parallelism. In brief overview, FIG. 5A illustrates embodiments of a multi-core system such as an appliance **200'** with n-cores, a total of cores numbers **1** through **N**. In one embodiment, work, load or network traffic can be distributed among a first core **505A**, a second core **505B**, a third core **505C**, a fourth core **505D**, a fifth core **505E**, a sixth core **505F**, a seventh core **505G**, and so on such that distribution is across all or two or more of the n cores **505N** (hereinafter referred to collectively as cores **505**.) There may be multiple VIPs **275** each running on a respective core of the plurality of cores. There may be multiple packet engines **240** each running on a respective core of the plurality of cores. Any of the approaches used may lead to different, varying or similar work load or performance level **515** across any of the cores. For a functional parallelism approach, each core may run a different function of the functionalities provided by the packet engine, a VIP **275** or appliance **200**. In a data parallelism approach, data may be paralleled or distributed across the cores based on the Network Interface Card (NIC) or VIP **275** receiving the data. In another data parallelism approach, processing may be distributed across the cores by distributing data flows to each core.

In further detail to FIG. 5A, in some embodiments, load, work or network traffic can be distributed among cores **505** according to functional parallelism **500**. Functional parallelism may be based on each core performing one or more respective functions. In some embodiments, a first core may perform a first function while a second core performs a second function. In functional parallelism approach, the functions to be performed by the multi-core system are divided and distributed to each core according to functionality. In some embodiments, functional parallelism may be referred to as task parallelism and may be achieved when each processor or core executes a different process or function on the same or different data. The core or processor may execute the same or different code. In some cases, different execution threads or code may communicate with one another as they work. Communication may take place to pass data from one thread to the next as part of a workflow.

In some embodiments, distributing work across the cores **505** according to functional parallelism **500**, can comprise distributing network traffic according to a particular function such as network input/output management (NW I/O) **510A**, secure sockets layer (SSL) encryption and decryption **510B** and transmission control protocol (TCP) functions **510C**.

42

This may lead to a work, performance or computing load **515** based on a volume or level of functionality being used. In some embodiments, distributing work across the cores **505** according to data parallelism **540**, can comprise distributing an amount of work **515** based on distributing data associated with a particular hardware or software component. In some embodiments, distributing work across the cores **505** according to flow-based data parallelism **520**, can comprise distributing data based on a context or flow such that the amount of work **515A-N** on each core may be similar, substantially equal or relatively evenly distributed.

In the case of the functional parallelism approach, each core may be configured to run one or more functionalities of the plurality of functionalities provided by the packet engine or VIP of the appliance. For example, core **1** may perform network I/O processing for the appliance **200'** while core **2** performs TCP connection management for the appliance. Likewise, core **3** may perform SSL offloading while core **4** may perform layer **7** or application layer processing and traffic management. Each of the cores may perform the same function or different functions. Each of the cores may perform more than one function. Any of the cores may run any of the functionality or portions thereof identified and/or described in conjunction with FIGS. 2A and 2B. In this the approach, the work across the cores may be divided by function in either a coarse-grained or fine-grained manner. In some cases, as illustrated in FIG. 5A, division by function may lead to different cores running at different levels of performance or load **515**.

In the case of the functional parallelism approach, each core may be configured to run one or more functionalities of the plurality of functionalities provided by the packet engine of the appliance. For example, core **1** may perform network I/O processing for the appliance **200'** while core **2** performs TCP connection management for the appliance. Likewise, core **3** may perform SSL offloading while core **4** may perform layer **7** or application layer processing and traffic management. Each of the cores may perform the same function or different functions. Each of the cores may perform more than one function. Any of the cores may run any of the functionality or portions thereof identified and/or described in conjunction with FIGS. 2A and 2B. In this the approach, the work across the cores may be divided by function in either a coarse-grained or fine-grained manner. In some cases, as illustrated in FIG. 5A division by function may lead to different cores running at different levels of load or performance.

The functionality or tasks may be distributed in any arrangement and scheme. For example, FIG. 5B illustrates a first core, Core **1 505A**, processing applications and processes associated with network I/O functionality **510A**. Network traffic associated with network I/O, in some embodiments, can be associated with a particular port number. Thus, outgoing and incoming packets having a port destination associated with NW I/O **510A** will be directed towards Core **1 505A** which is dedicated to handling all network traffic associated with the NW I/O port. Similarly, Core **2 505B** is dedicated to handling functionality associated with SSL processing and Core **4 505D** may be dedicated handling all TCP level processing and functionality.

While FIG. 5A illustrates functions such as network I/O, SSL and TCP, other functions can be assigned to cores. These other functions can include any one or more of the functions or operations described herein. For example, any of the functions described in conjunction with FIGS. 2A and 2B may be distributed across the cores on a functionality basis. In some cases, a first VIP **275A** may run on a first core while a second VIP **275B** with a different configuration may run on a second

core. In some embodiments, each core **505** can handle a particular functionality such that each core **505** can handle the processing associated with that particular function. For example, Core **2 505B** may handle SSL offloading while Core **4 505D** may handle application layer processing and traffic management.

In other embodiments, work, load or network traffic may be distributed among cores **505** according to any type and form of data parallelism **540**. In some embodiments, data parallelism may be achieved in a multi-core system by each core performing the same task or functionally on different pieces of distributed data. In some embodiments, a single execution thread or code controls operations on all pieces of data. In other embodiments, different threads or instructions control the operation, but may execute the same code. In some embodiments, data parallelism is achieved from the perspective of a packet engine, vServers (VIPs) **275A-C**, network interface cards (NIC) **542D-E** and/or any other networking hardware or software included on or associated with an appliance **200**. For example, each core may run the same packet engine or VIP code or configuration but operate on different sets of distributed data. Each networking hardware or software construct can receive different, varying or substantially the same amount of data, and as a result may have varying, different or relatively the same amount of load **515**.

In the case of a data parallelism approach, the work may be divided up and distributed based on VIPs, NICs and/or data flows of the VIPs or NICs. In one of these approaches, the work of the multi-core system may be divided or distributed among the VIPs by having each VIP work on a distributed set of data. For example, each core may be configured to run one or more VIPs. Network traffic may be distributed to the core for each VIP handling that traffic. In another of these approaches, the work of the appliance may be divided or distributed among the cores based on which NIC receives the network traffic. For example, network traffic of a first NIC may be distributed to a first core while network traffic of a second NIC may be distributed to a second core. In some cases, a core may process data from multiple NICs.

While FIG. **5A** illustrates a single vServer associated with a single core **505**, as is the case for VIP1 **275A**, VIP2 **275B** and VIP3 **275C**. In some embodiments, a single vServer can be associated with one or more cores **505**. In contrast, one or more vServers can be associated with a single core **505**. Associating a vServer with a core **505** may include that core **505** to process all functions associated with that particular vServer. In some embodiments, each core executes a VIP having the same code and configuration. In other embodiments, each core executes a VIP having the same code but different configuration. In some embodiments, each core executes a VIP having different code and the same or different configuration.

Like vServers, NICs can also be associated with particular cores **505**. In many embodiments, NICs can be connected to one or more cores **505** such that when a NIC receives or transmits data packets, a particular core **505** handles the processing involved with receiving and transmitting the data packets. In one embodiment, a single NIC can be associated with a single core **505**, as is the case with NIC1 **542D** and NIC2 **542E**. In other embodiments, one or more NICs can be associated with a single core **505**. In other embodiments, a single NIC can be associated with one or more cores **505**. In these embodiments, load could be distributed amongst the one or more cores **505** such that each core **505** processes a substantially similar amount of load. A core **505** associated with a NIC may process all functions and/or data associated with that particular NIC.

While distributing work across cores based on data of VIPs or NICs may have a level of independency, in some embodiments, this may lead to unbalanced use of cores as illustrated by the varying loads **515** of FIG. **5A**.

In some embodiments, load, work or network traffic can be distributed among cores **505** based on any type and form of data flow. In another of these approaches, the work may be divided or distributed among cores based on data flows. For example, network traffic between a client and a server traversing the appliance may be distributed to and processed by one core of the plurality of cores. In some cases, the core initially establishing the session or connection may be the core for which network traffic for that session or connection is distributed. In some embodiments, the data flow is based on any unit or portion of network traffic, such as a transaction, a request/response communication or traffic originating from an application on a client. In this manner and in some embodiments, data flows between clients and servers traversing the appliance **200'** may be distributed in a more balanced manner than the other approaches.

In flow-based data parallelism **520**, distribution of data is related to any type of flow of data, such as request/response pairings, transactions, sessions, connections or application communications. For example, network traffic between a client and a server traversing the appliance may be distributed to and processed by one core of the plurality of cores. In some cases, the core initially establishing the session or connection may be the core for which network traffic for that session or connection is distributed. The distribution of data flow may be such that each core **505** carries a substantially equal or relatively evenly distributed amount of load, data or network traffic.

In some embodiments, the data flow is based on any unit or portion of network traffic, such as a transaction, a request/response communication or traffic originating from an application on a client. In this manner and in some embodiments, data flows between clients and servers traversing the appliance **200'** may be distributed in a more balanced manner than the other approached. In one embodiment, data flow can be distributed based on a transaction or a series of transactions. This transaction, in some embodiments, can be between a client and a server and can be characterized by an IP address or other packet identifier. For example, Core **1 505A** can be dedicated to transactions between a particular client and a particular server, therefore the load **515A** on Core **1 505A** may be comprised of the network traffic associated with the transactions between the particular client and server. Allocating the network traffic to Core **1 505A** can be accomplished by routing all data packets originating from either the particular client or server to Core **1 505A**.

While work or load can be distributed to the cores based in part on transactions, in other embodiments load or work can be allocated on a per packet basis. In these embodiments, the appliance **200** can intercept data packets and allocate them to a core **505** having the least amount of load. For example, the appliance **200** could allocate a first incoming data packet to Core **1 505A** because the load **515A** on Core **1** is less than the load **515B-N** on the rest of the cores **505B-N**. Once the first data packet is allocated to Core **1 505A**, the amount of load **515A** on Core **1 505A** is increased proportional to the amount of processing resources needed to process the first data packet. When the appliance **200** intercepts a second data packet, the appliance **200** will allocate the load to Core **4 505D** because Core **4 505D** has the second least amount of load. Allocating data packets to the core with the least amount

of load can, in some embodiments, ensure that the load **515A-N** distributed to each core **505** remains substantially equal.

In other embodiments, load can be allocated on a per unit basis where a section of network traffic is allocated to a particular core **505**. The above-mentioned example illustrates load balancing on a per/packet basis. In other embodiments, load can be allocated based on a number of packets such that every 10, 100 or 1000 packets are allocated to the core **505** having the least amount of load. The number of packets allocated to a core **505** can be a number determined by an application, user or administrator and can be any number greater than zero. In still other embodiments, load can be allocated based on a time metric such that packets are distributed to a particular core **505** for a predetermined amount of time. In these embodiments, packets can be distributed to a particular core **505** for five milliseconds or for any period of time determined by a user, program, system, administrator or otherwise. After the predetermined time period elapses, data packets are transmitted to a different core **505** for the predetermined period of time.

Flow-based data parallelism methods for distributing work, load or network traffic among the one or more cores **505** can comprise any combination of the above-mentioned embodiments. These methods can be carried out by any part of the appliance **200**, by an application or set of executable instructions executing on one of the cores **505**, such as the packet engine, or by any application, program or agent executing on a computing device in communication with the appliance **200**.

The functional and data parallelism computing schemes illustrated in FIG. **5A** can be combined in any manner to generate a hybrid parallelism or distributed processing scheme that encompasses function parallelism **500**, data parallelism **540**, flow-based data parallelism **520** or any portions thereof. In some cases, the multi-core system may use any type and form of load balancing schemes to distribute load among the one or more cores **505**. The load balancing scheme may be used in any combination with any of the functional and data parallelism schemes or combinations thereof.

Illustrated in FIG. **5B** is an embodiment of a multi-core system **545**, which may be any type and form of one or more systems, appliances, devices or components. This system **545**, in some embodiments, can be included within an appliance **200** having one or more processing cores **505A-N**. The system **545** can further include one or more packet engines (PE) or packet processing engines (PPE) **548A-N** communicating with a memory bus **556**. The memory bus may be used to communicate with the one or more processing cores **505A-N**. Also included within the system **545** can be one or more network interface cards (NIC) **552** and a flow distributor **550** which can further communicate with the one or more processing cores **505A-N**. The flow distributor **550** can comprise a Receive Side Scaler (RSS) or Receive Side Scaling (RSS) module **560**.

Further referring to FIG. **5B**, and in more detail, in one embodiment the packet engine(s) **548A-N** can comprise any portion of the appliance **200** described herein, such as any portion of the appliance described in FIGS. **2A** and **2B**. The packet engine(s) **548A-N** can, in some embodiments, comprise any of the following elements: the packet engine **240**, a network stack **267**; a cache manager **232**; a policy engine **236**; a compression engine **238**; an encryption engine **234**; a GUI **210**; a CLI **212**; shell services **214**; monitoring programs **216**; and any other software or hardware element able to receive data packets from one of either the memory bus **556** or the one of more cores **505A-N**. In some embodiments, the packet

engine(s) **548A-N** can comprise one or more vServers **275A-N**, or any portion thereof. In other embodiments, the packet engine(s) **548A-N** can provide any combination of the following functionalities: SSL VPN **280**; Intranet UP **282**; switching **284**; DNS **286**; packet acceleration **288**; App FW **280**; monitoring such as the monitoring provided by a monitoring agent **197**; functionalities associated with functioning as a TCP stack; load balancing; SSL offloading and processing; content switching; policy evaluation; caching; compression; encoding; decompression; decoding; application firewall functionalities; XML processing and acceleration; and SSL VPN connectivity.

The packet engine(s) **548A-N** can, in some embodiments, be associated with a particular server, user, client or network. When a packet engine **548** becomes associated with a particular entity, that packet engine **548** can process data packets associated with that entity. For example, should a packet engine **548** be associated with a first user, that packet engine **548** will process and operate on packets generated by the first user, or packets having a destination address associated with the first user. Similarly, the packet engine **548** may choose not to be associated with a particular entity such that the packet engine **548** can process and otherwise operate on any data packets not generated by that entity or destined for that entity.

In some instances, the packet engine(s) **548A-N** can be configured to carry out any of the functional and/or data parallelism schemes illustrated in FIG. **5A**. In these instances, the packet engine(s) **548A-N** can distribute functions or data among the processing cores **505A-N** so that the distribution is according to the parallelism or distribution scheme. In some embodiments, a single packet engine(s) **548A-N** carries out a load balancing scheme, while in other embodiments one or more packet engine(s) **548A-N** carry out a load balancing scheme. Each core **505A-N**, in one embodiment, can be associated with a particular packet engine **548** such that load balancing can be carried out by the packet engine. Load balancing may in this embodiment, require that each packet engine **548A-N** associated with a core **505** communicate with the other packet engines associated with cores so that the packet engines **548A-N** can collectively determine where to distribute load. One embodiment of this process can include an arbiter that receives votes from each packet engine for load. The arbiter can distribute load to each packet engine **548A-N** based in part on the age of the engine's vote and in some cases a priority value associated with the current amount of load on an engine's associated core **505**.

Any of the packet engines running on the cores may run in user mode, kernel or any combination thereof. In some embodiments, the packet engine operates as an application or program running in user or application space. In these embodiments, the packet engine may use any type and form of interface to access any functionality provided by the kernel. In some embodiments, the packet engine operates in kernel mode or as part of the kernel. In some embodiments, a first portion of the packet engine operates in user mode while a second portion of the packet engine operates in kernel mode. In some embodiments, a first packet engine on a first core executes in kernel mode while a second packet engine on a second core executes in user mode. In some embodiments, the packet engine or any portions thereof operates on or in conjunction with the NIC or any drivers thereof.

In some embodiments the memory bus **556** can be any type and form of memory or computer bus. While a single memory bus **556** is depicted in FIG. **5B**, the system **545** can comprise any number of memory buses **556**. In one embodiment, each packet engine **548** can be associated with one or more individual memory buses **556**.

47

The NIC 552 can in some embodiments be any of the network interface cards or mechanisms described herein. The NIC 552 can have any number of ports. The NIC can be designed and constructed to connect to any type and form of network 104. While a single NIC 552 is illustrated, the system 545 can comprise any number of NICs 552. In some embodiments, each core 505A-N can be associated with one or more single NICs 552. Thus, each core 505 can be associated with a single NIC 552 dedicated to a particular core 505. The cores 505A-N can comprise any of the processors described herein. Further, the cores 505A-N can be configured according to any of the core 505 configurations described herein. Still further, the cores 505A-N can have any of the core 505 functionalities described herein. While FIG. 5B illustrates seven cores 505A-G, any number of cores 505 can be included within the system 545. In particular, the system 545 can comprise "N" cores, where "N" is a whole number greater than zero.

A core may have or use memory that is allocated or assigned for use to that core. The memory may be considered private or local memory of that core and only accessible by that core. A core may have or use memory that is shared or assigned to multiple cores. The memory may be considered public or shared memory that is accessible by more than one core. A core may use any combination of private and public memory. With separate address spaces for each core, some level of coordination is eliminated from the case of using the same address space. With a separate address space, a core can perform work on information and data in the core's own address space without worrying about conflicts with other cores. Each packet engine may have a separate memory pool for TCP and/or SSL connections.

Further referring to FIG. 5B, any of the functionality and/or embodiments of the cores 505 described above in connection with FIG. 5A can be deployed in any embodiment of the virtualized environment described above in connection with FIGS. 4A and 4B. Instead of the functionality of the cores 505 being deployed in the form of a physical processor 505, such functionality may be deployed in a virtualized environment on any computing device 100, such as a client 102, server 106 or appliance 200. In other embodiments, instead of the functionality of the cores 505 being deployed in the form of an appliance or a single device, the functionality may be deployed across multiple devices in any arrangement. For example, one device may comprise two or more cores and another device may comprise two or more cores. For example, a multi-core system may include a cluster of computing devices, a server farm or network of computing devices. In some embodiments, instead of the functionality of the cores 505 being deployed in the form of cores, the functionality may be deployed on a plurality of processors, such as a plurality of single core processors.

In one embodiment, the cores 505 may be any type and form of processor. In some embodiments, a core can function substantially similar to any processor or central processing unit described herein. In some embodiment, the cores 505 may comprise any portion of any processor described herein. While FIG. 5A illustrates seven cores, there can exist any "N" number of cores within an appliance 200, where "N" is any whole number greater than one. In some embodiments, the cores 505 can be installed within a common appliance 200, while in other embodiments the cores 505 can be installed within one or more appliance(s) 200 communicatively connected to one another. The cores 505 can in some embodiments comprise graphics processing software, while in other embodiments the cores 505 provide general processing capabilities. The cores 505 can be installed physically near each other and/or can be communicatively connected to each other.

48

The cores may be connected by any type and form of bus or subsystem physically and/or communicatively coupled to the cores for transferring data between to, from and/or between the cores.

While each core 505 can comprise software for communicating with other cores, in some embodiments a core manager (not shown) can facilitate communication between each core 505. In some embodiments, the kernel may provide core management. The cores may interface or communicate with each other using a variety of interface mechanisms. In some embodiments, core to core messaging may be used to communicate between cores, such as a first core sending a message or data to a second core via a bus or subsystem connecting the cores. In some embodiments, cores may communicate via any type and form of shared memory interface. In one embodiment, there may be one or more memory locations shared among all the cores. In some embodiments, each core may have separate memory locations shared with each other core. For example, a first core may have a first shared memory with a second core and a second share memory with a third core. In some embodiments, cores may communicate via any type of programming or API, such as function calls via the kernel. In some embodiments, the operating system may recognize and support multiple core devices and provide interfaces and API for inter-core communications.

The flow distributor 550 can be any application, program, library, script, task, service, process or any type and form of executable instructions executing on any type and form of hardware. In some embodiments, the flow distributor 550 may any design and construction of circuitry to perform any of the operations and functions described herein. In some embodiments, the flow distributor distribute, forwards, routes, controls and/or manage the distribution of data packets among the cores 505 and/or packet engine or VIPs running on the cores. The flow distributor 550, in some embodiments, can be referred to as an interface master. In one embodiment, the flow distributor 550 comprises a set of executable instructions executing on a core or processor of the appliance 200. In another embodiment, the flow distributor 550 comprises a set of executable instructions executing on a computing machine in communication with the appliance 200. In some embodiments, the flow distributor 550 comprises a set of executable instructions executing on a NIC, such as firmware. In still other embodiments, the flow distributor 550 comprises any combination of software and hardware to distribute data packets among cores or processors. In one embodiment, the flow distributor 550 executes on at least one of the cores 505A-N, while in other embodiments a separate flow distributor 550 assigned to each core 505A-N executes on an associated core 505A-N. The flow distributor may use any type and form of statistical or probabilistic algorithms or decision making to balance the flows across the cores. The hardware of the appliance, such as a NIC, or the kernel may be designed and constructed to support sequential operations across the NICs and/or cores.

In embodiments where the system 545 comprises one or more flow distributors 550, each flow distributor 550 can be associated with a processor 505 or a packet engine 548. The flow distributors 550 can comprise an interface mechanism that allows each flow distributor 550 to communicate with the other flow distributors 550 executing within the system 545. In one instance, the one or more flow distributors 550 can determine how to balance load by communicating with each other. This process can operate substantially similarly to the process described above for submitting votes to an arbiter which then determines which flow distributor 550 should receive the load. In other embodiments, a first flow distributor

550' can identify the load on an associated core and determine whether to forward a first data packet to the associated core based on any of the following criteria: the load on the associated core is above a predetermined threshold; the load on the associated core is below a predetermined threshold; the load on the associated core is less than the load on the other cores; or any other metric that can be used to determine where to forward data packets based in part on the amount of load on a processor.

The flow distributor **550** can distribute network traffic among the cores **505** according to a distribution, computing or load balancing scheme such as those described herein. In one embodiment, the flow distributor can distribute network traffic according to any one of a functional parallelism distribution scheme **550**, a data parallelism load distribution scheme **540**, a flow-based data parallelism distribution scheme **520**, or any combination of these distribution scheme or any load balancing scheme for distributing load among multiple processors. The flow distributor **550** can therefore act as a load distributor by taking in data packets and distributing them across the processors according to an operative load balancing or distribution scheme. In one embodiment, the flow distributor **550** can comprise one or more operations, functions or logic to determine how to distribute packers, work or load accordingly. In still other embodiments, the flow distributor **550** can comprise one or more sub operations, functions or logic that can identify a source address and a destination address associated with a data packet, and distribute packets accordingly.

In some embodiments, the flow distributor **550** can comprise a receive-side scaling (RSS) network driver, module **560** or any type and form of executable instructions which distribute data packets among the one or more cores **505**. The RSS module **560** can comprise any combination of hardware and software. In some embodiments, the RSS module **560** works in conjunction with the flow distributor **550** to distribute data packets across the cores **505A-N** or among multiple processors in a multi-processor network. The RSS module **560** can execute within the NIC **552** in some embodiments, and in other embodiments can execute on any one of the cores **505**.

In some embodiments, the RSS module **560** uses the MICROSOFT receive-side-scaling (RSS) scheme. In one embodiment, RSS is a Microsoft Scalable Networking initiative technology that enables receive processing to be balanced across multiple processors in the system while maintaining in-order delivery of the data. The RSS may use any type and form of hashing scheme to determine a core or processor for processing a network packet.

The RSS module **560** can apply any type and form hash function such as the Toeplitz hash function. The hash function may be applied to the hash type or any the sequence of values. The hash function may be a secure hash of any security level or is otherwise cryptographically secure. The hash function may use a hash key. The size of the key is dependent upon the hash function. For the Toeplitz hash, the size may be 40 bytes for IPv6 and 16 bytes for IPv4.

The hash function may be designed and constructed based on any one or more criteria or design goals. In some embodiments, a hash function may be used that provides an even distribution of hash result for different hash inputs and different hash types, including TCP/IPv4, TCP/IPv6, IPv4, and IPv6 headers. In some embodiments, a hash function may be used that provides a hash result that is evenly distributed when a small number of buckets are present (for example, two or four). In some embodiments, hash function may be used that provides a hash result that is randomly distributed when a

large number of buckets were present (for example, 64 buckets). In some embodiments, the hash function is determined based on a level of computational or resource usage. In some embodiments, the hash function is determined based on ease or difficulty of implementing the hash in hardware. In some embodiments, the hash function is determined based on the ease or difficulty of a malicious remote host to send packets that would all hash to the same bucket.

The RSS may generate hashes from any type and form of input, such as a sequence of values. This sequence of values can include any portion of the network packet, such as any header, field or payload of network packet, or portions thereof. In some embodiments, the input to the hash may be referred to as a hash type and include any tuples of information associated with a network packet or data flow, such as any of the following: a four tuple comprising at least two IP addresses and two ports; a four tuple comprising any four sets of values; a six tuple; a two tuple; and/or any other sequence of numbers or values. The following are example of hash types that may be used by RSS:

4-tuple of source TCP Port, source IP version 4 (IPv4) address, destination TCP Port, and destination IPv4 address.

4-tuple of source TCP Port, source IP version 6 (IPv6) address, destination TCP Port, and destination IPv6 address.

2-tuple of source IPv4 address, and destination IPv4 address.

2-tuple of source IPv6 address, and destination IPv6 address.

2-tuple of source IPv6 address, and destination IPv6 address, including support for parsing IPv6 extension headers.

The hash result or any portion thereof may used to identify a core or entity, such as a packet engine or VIP, for distributing a network packet. In some embodiments, one or more hash bits or mask are applied to the hash result. The hash bit or mask may be any number of bits or bytes. A NIC may support any number of bits, such as seven bits. The network stack may set the actual number of bits to be used during initialization. The number will be between 1 and 7, inclusive.

The hash result may be used to identify the core or entity via any type and form of table, such as a bucket table or indirection table. In some embodiments, the number of hash-result bits are used to index into the table. The range of the hash mask may effectively define the size of the indirection table. Any portion of the hash result or the hash result itself may be used to index the indirection table. The values in the table may identify any of the cores or processor, such as by a core or processor identifier. In some embodiments, all of the cores of the multi-core system are identified in the table. In other embodiments, a port of the cores of the multi-core system are identified in the table. The indirection table may comprise any number of buckets for example 2 to 128 buckets that may be indexed by a hash mask. Each bucket may comprise a range of index values that identify a core or processor. In some embodiments, the flow controller and/or RSS module may rebalance the network load by changing the indirection table.

In some embodiments, the multi-core system **575** does not include a RSS driver or RSS module **560**. In some of these embodiments, a software steering module (not shown) or a software embodiment of the RSS module within the system can operate in conjunction with or as part of the flow distributor **550** to steer packets to cores **505** within the multi-core system **575**.

51

The flow distributor **550**, in some embodiments, executes within any module or program on the appliance **200**, on any one of the cores **505** and on any one of the devices or components included within the multi-core system **575**. In some embodiments, the flow distributor **550'** can execute on the first core **505A**, while in other embodiments the flow distributor **550'** can execute on the NIC **552**. In still other embodiments, an instance of the flow distributor **550'** can execute on each core **505** included in the multi-core system **575**. In this embodiment, each instance of the flow distributor **550'** can communicate with other instances of the flow distributor **550'** to forward packets back and forth across the cores **505**. There exist situations where a response to a request packet may not be processed by the same core, i.e. the first core processes the request while the second core processes the response. In these situations, the instances of the flow distributor **550'** can intercept the packet and forward it to the desired or correct core **505**, i.e. a flow distributor instance **550'** can forward the response to the first core. Multiple instances of the flow distributor **550'** can execute on any number of cores **505** and any combination of cores **505**.

The flow distributor may operate responsive to any one or more rules or policies. The rules may identify a core or packet processing engine to receive a network packet, data or data flow. The rules may identify any type and form of tuple information related to a network packet, such as a 4-tuple of source and destination IP address and source and destination ports. Based on a received packet matching the tuple specified by the rule, the flow distributor may forward the packet to a core or packet engine. In some embodiments, the packet is forwarded to a core via shared memory and/or core to core messaging.

Although FIG. **5B** illustrates the flow distributor **550** as executing within the multi-core system **575**, in some embodiments the flow distributor **550** can execute on a computing device or appliance remotely located from the multi-core system **575**. In such an embodiment, the flow distributor **550** can communicate with the multi-core system **575** to take in data packets and distribute the packets across the one or more cores **505**. The flow distributor **550** can, in one embodiment, receive data packets destined for the appliance **200**, apply a distribution scheme to the received data packets and distribute the data packets to the one or more cores **505** of the multi-core system **575**. In one embodiment, the flow distributor **550** can be included in a router or other appliance such that the router can target particular cores **505** by altering meta data associated with each packet so that each packet is targeted towards a sub-node of the multi-core system **575**. In such an embodiment, CISCO's vn-tag mechanism can be used to alter or tag each packet with the appropriate meta data.

Illustrated in FIG. **5C** is an embodiment of a multi-core system **575** comprising one or more processing cores **505A-N**. In brief overview, one of the cores **505** can be designated as a control core **505A** and can be used as a control plane **570** for the other cores **505**. The other cores may be secondary cores which operate in a data plane while the control core provides the control plane. The cores **505A-N** may share a global cache **580**. While the control core provides a control plane, the other cores in the multi-core system form or provide a data plane. These cores perform data processing functionality on network traffic while the control provides initialization, configuration and control of the multi-core system.

Further referring to FIG. **5C**, and in more detail, the cores **505A-N** as well as the control core **505A** can be any processor described herein. Furthermore, the cores **505A-N** and the control core **505A** can be any processor able to function within the system **575** described in FIG. **5C**. Still further, the

52

cores **505A-N** and the control core **505A** can be any core or group of cores described herein. The control core may be a different type of core or processor than the other cores. In some embodiments, the control may operate a different packet engine or have a packet engine configured differently than the packet engines of the other cores.

Any portion of the memory of each of the cores may be allocated to or used for a global cache that is shared by the cores. In brief overview, a predetermined percentage or predetermined amount of each of the memory of each core may be used for the global cache. For example, 50% of each memory of each code may be dedicated or allocated to the shared global cache. That is, in the illustrated embodiment, 2 GB of each core excluding the control plane core or core **1** may be used to form a 28 GB shared global cache. The configuration of the control plane such as via the configuration services may determine the amount of memory used for the shared global cache. In some embodiments, each core may provide a different amount of memory for use by the global cache. In other embodiments, any one core may not provide any memory or use the global cache. In some embodiments, any of the cores may also have a local cache in memory not allocated to the global shared memory. Each of the cores may store any portion of network traffic to the global shared cache. Each of the cores may check the cache for any content to use in a request or response. Any of the cores may obtain content from the global shared cache to use in a data flow, request or response.

The global cache **580** can be any type and form of memory or storage element, such as any memory or storage element described herein. In some embodiments, the cores **505** may have access to a predetermined amount of memory (i.e. 32 GB or any other memory amount commensurate with the system **575**). The global cache **580** can be allocated from that predetermined amount of memory while the rest of the available memory can be allocated among the cores **505**. In other embodiments, each core **505** can have a predetermined amount of memory. The global cache **580** can comprise an amount of the memory allocated to each core **505**. This memory amount can be measured in bytes, or can be measured as a percentage of the memory allocated to each core **505**. Thus, the global cache **580** can comprise 1 GB of memory from the memory associated with each core **505**, or can comprise 20 percent or one-half of the memory associated with each core **505**. In some embodiments, only a portion of the cores **505** provide memory to the global cache **580**, while in other embodiments the global cache **580** can comprise memory not allocated to the cores **505**.

Each core **505** can use the global cache **580** to store network traffic or cache data. In some embodiments, the packet engines of the core use the global cache to cache and use data stored by the plurality of packet engines. For example, the cache manager of FIG. **2A** and cache functionality of FIG. **2B** may use the global cache to share data for acceleration. For example, each of the packet engines may store responses, such as HTML data, to the global cache. Any of the cache managers operating on a core may access the global cache to server caches responses to client requests.

In some embodiments, the cores **505** can use the global cache **580** to store a port allocation table which can be used to determine data flow based in part on ports. In other embodiments, the cores **505** can use the global cache **580** to store an address lookup table or any other table or list that can be used by the flow distributor to determine where to direct incoming and outgoing data packets. The cores **505** can, in some embodiments read from and write to cache **580**, while in other

53

embodiments the cores **505** can only read from or write to cache **580**. The cores may use the global cache to perform core to core communications.

The global cache **580** may be sectioned into individual memory sections where each section can be dedicated to a particular core **505**. In one embodiment, the control core **505A** can receive a greater amount of available cache, while the other cores **505** can receiving varying amounts or access to the global cache **580**.

In some embodiments, the system **575** can comprise a control core **505A**. While FIG. **5C** illustrates core **1 505A** as the control core, the control core can be any core within the appliance **200** or multi-core system. Further, while only a single control core is depicted, the system **575** can comprise one or more control cores each having a level of control over the system. In some embodiments, one or more control cores can each control a particular aspect of the system **575**. For example, one core can control deciding which distribution scheme to use, while another core can determine the size of the global cache **580**.

The control plane of the multi-core system may be the designation and configuration of a core as the dedicated management core or as a master core. This control plane core may provide control, management and coordination of operation and functionality the plurality of cores in the multi-core system. This control plane core may provide control, management and coordination of allocation and use of memory of the system among the plurality of cores in the multi-core system, including initialization and configuration of the same. In some embodiments, the control plane includes the flow distributor for controlling the assignment of data flows to cores and the distribution of network packets to cores based on data flows. In some embodiments, the control plane core runs a packet engine and in other embodiments, the control plane core is dedicated to management and control of the other cores of the system.

The control core **505A** can exercise a level of control over the other cores **505** such as determining how much memory should be allocated to each core **505** or determining which core **505** should be assigned to handle a particular function or hardware/software entity. The control core **505A**, in some embodiments, can exercise control over those cores **505** within the control plan **570**. Thus, there can exist processors outside of the control plane **570** which are not controlled by the control core **505A**. Determining the boundaries of the control plane **570** can include maintaining, by the control core **505A** or agent executing within the system **575**, a list of those cores **505** controlled by the control core **505A**. The control core **505A** can control any of the following: initialization of a core; determining when a core is unavailable; re-distributing load to other cores **505** when one core fails; determining which distribution scheme to implement; determining which core should receive network traffic; determining how much cache should be allocated to each core; determining whether to assign a particular function or element to a particular core; determining whether to permit cores to communicate with one another; determining the size of the global cache **580**; and any other determination of a function, configuration or operation of the cores within the system **575**.

F. Systems and Methods for Aggregating Multi-Core Performance and Trace Data

The systems and methods described in this section are directed towards aggregating performance and trace data of a multi-core, multi-packet-engine appliance which manages network traffic. In various embodiments, the appliance is disposed in a network intermediary to a plurality of clients and a plurality of servers which can provide network services

54

or web content to the clients. The multi-core appliance can manage connections between plural clients and plural servers as well as manage web content exchanged between clients and servers.

In overview and referring now to FIG. **6A**, an embodiment of a system **600** for aggregating performance or trace data is depicted in block diagram form. The system **600** can be configured for operation on an appliance **200**. The appliance **200**, as described elsewhere herein, can be adapted to operate with the system **600** and provide information about network performance and trace data of the appliance. In various embodiments, the performance and trace aggregating system **600** comprises packet engines **548A-N**, shared memory allocations **122B-N**, an aggregator **610**, and an agent interface **615**. Each packet engine can manage network traffic between one or more clients **102** and one or more servers **106**. The aggregator **610** can communicate with the packet engines **548A-N** via shared memory allocations **122B-N**, as well as collect data through the shared memory allocations. The aggregator **610** can further communicate with agents **620A**, **620B**, which can request and receive information about the appliance's performance and trace data through the agent interface **615**. The agent interface **615** can function as an intermediary between the agents **620** and the aggregator **610**.

In various embodiments and in further detail, the packet engines **548A-N** execute on a plurality of cores **505A-N** of the appliance **200**. Each packet engine can execute on a respective core and manage network traffic, e.g., exchange of web content, between one or more clients **102** and one or more servers **106** in communication with the appliance **200**. Aspects of the packet engines **548A-N** and cores **505A-N** are described elsewhere herein. Web content can include, without being limited to, HTTP or HTTPS requests, HTTP or HTTPS responses, electronic mail, web pages, video data, music or audio data, web forms, statistical data, plots, maps, books, advertisements, news or journal text, pictures, compiled data descriptive of an item, event, weather, individual, group, organization, business, or geographical area. Each packet engine **548A-N** can manage services provided by servers **106**. Web services can include, without being limited to, providing HTTP or HTTPS responses in reply to a request, searching for specified information, carrying out calculations, compiling data, compiling statistics, providing for purchasing or sale of an item, executing a computer game, providing for electronic communication with another networked entity, etc. In various embodiments, each packet engine **548A-N** communicates with an allocation **122B-N** in memory **122**.

In certain embodiments, each packet engine **548A-N** is adapted to maintain one or more records **605PA-N**, **605TA-N** (generally **605**) of certain aspects of its operation, e.g., certain statistics or data about its operation as a virtual server (VIP **275A-N**), data about servers to which each packet engine is currently maintaining active connections or has had connections to, data about a service managed by each packet engine, etc. In some embodiments, each packet engine maintains separate records for performance data **650P** and trace data **605T**. In some embodiments, each packet engine maintains separate records for types of performance data, e.g., a record for VIP data, a record for subscribed service data. In some embodiments, each packet engine maintains combined data records, e.g., one record for performance and trace data. In various embodiments, the records **605** provide information about performance statistics of each packet engine, and/or contain trace data captured by each packet engine. Each record **605A-N** can comprise a data structure containing multiple records and types of data. Each record can be maintained

with each core in memory **122**, cache **140**, in a data buffer, or any type of data storage element accessible to each packet engine or its core.

Performance data **605P** of the VIP **275** aspects of operation and/or performance data for each packet engine can include, without being limited to, any of the following elements: the IP address and port configuration of the VIP, the state of the VIP, e.g., up, down, out of service, an indication of load balancing configuration of the VIP, number of hits on the VIP requesting web content, a measure of traffic passed through the VIP, e.g., kbps, Mbps, Gbps, the persistence configuration for the VIP, a measure of the number of times an error page was generated by the VIP, a measure of the volume of traffic passing through the VIP, e.g., number of packets passing through the VIP, average size of a packet, number of active services that are bound to the VIP, the number of client connections to the VIP, the rate at which client connections are added or deleted from the VIP, the number of server connections from the VIP in an "open" state, the number of current server connections from the VIP, and any combination thereof.

In some embodiments, the records **605** are maintained by each packet engine during operation, e.g., each packet engine can record the information periodically at selected time intervals. In some instances, each packet engine periodically overwrites data when recording the information, e.g., the data is written to the same space in a record at periodic intervals, new data overwriting old data. In some instances, the data is written to memory in a manner to provide historical data, e.g., the data is recorded to a circular buffer having M sections. After M time intervals, older data is overwritten, however the buffer retains a historical record of M time intervals at any instant. In various embodiments, each packet engine can record the information stored in its respective record **605** or a portion thereof to a shared memory allocation **122B-N** or to non-shared memory. In some embodiments, the record is generated by the packet engine after it receives a request for information. A record generated after request can include information about all VIP operational aspects reportable by the packet engine, or can include information about operational aspects specified in the request.

Information about servers or service of a server for which the packet engine is managing service can include without being limited to, any of the following elements, which may be included in the performance data **605P**: the IP address and port configuration of the server, the state of the service provided by the server, e.g., up, down, out of service, the number of hits directed to the service, the number of hits directed to the service due to configured server persistence, the number of active connections to the service, a measure of traffic to the service, e.g., kbps, Mbps, Gbps, a bandwidth limit for the service, an average response time of the service, a measure of traffic volume to the service, e.g., number of packets, average size of packets, a weight index used in load balancing for the service, a running hits count which can be used in a round-robin load balancing algorithm, a number of connections to the service, a rate at which connections to the service are added or deleted, a maximum number of connections to the service at one time that has been reached, a number of connections to the service in an "open" state, e.g., an open and active state, a number of connections to the service residing in a reuse pool, e.g., open but not active, a number of connections to the service waiting in a serge queue, and any combination thereof. In some embodiments, each record **605** of server or service information is maintained by each packet engine during operation, e.g., each packet engine can record the information periodically at selected time intervals. In some instances, each packet engine periodically overwrites

data when recording the information, e.g., the data is written to the same space in the record at periodic intervals, new data overwriting old data. In some instances, the data is written to memory in a manner to provide historical data, e.g., the data is recorded to a circular buffer having M sections. After M time intervals, older data is overwritten, however the buffer retains a historical record of M time intervals at any instant. In various embodiments, each packet engine can record the information from the records **605** to a shared memory allocation **122B-N** or to non-shared memory. In some embodiments, each record **605** is generated by the packet engine after it receives a request for information. A record generated after a request can include all service information reportable by the packet engine, or can include service information specified in the request.

In various embodiments, each packet engine **548A-N** is adapted to capture trace data. Trace data can identify network traffic or any portion thereof flowing through the packet engine during a selected time interval. Trace data can include any information related to the receipt, transmittal and processing of network packets and any hardware and/or software related thereto, such as the packet engine, network stack and NIC. Trace data can include, without being limited to, any of the elements from the following list: requests for content, responses to requests for content, encrypted content, non-encrypted content, and any combination thereof. Trace data can include identification of headers, fields and/or payload of a packet. Trace data can include identification of a protocol of packet. Trace data can include identification of a layer of the network stack of a packet. Trace data can include identification of a NIC. Trace data include any temporal information related to receipt, transmittal or processing of a network packet, such any granularity of timestamp. Trace data include any information identifying operations, functions or actions of the packet engine in processing a network packet. Trace data may be represented in any format, arrangement, structure or representation.

In certain embodiments, each packet engine is further adapted to filter web content passed by the packet engine in assembling trace data. As an example, each packet engine can select for trace data web content in accordance with a filtering criterion selected from the following list: source IP address, source port number, destination IP address, destination port number, service name, virtual server name, state of service, and any combination thereof. Combinations of these elements can include boolean combinations and compound expressions, e.g., sourceip=105.2.1.2 AND servicename=map/boston. In various embodiments, trace data is generated after the packet engine receives a request for trace data, the request specifying information desired in the trace and a time interval during which the trace is to be taken. If no information is specified, all available trace data can be reported for a current selected time interval. Trace data can be stored in separate trace records **605TA-N** for each packet engine.

Memory allocations **122B-N** can be defined in memory **122**. Details of memory **122** are described elsewhere herein. Memory allocations **122B-N** can comprise locations and space in memory reserved for use by each packet engine and for use by the aggregator **610**. The memory allocations **122B-N** can be defined by each packet engine **548A-N**, e.g., by using an IOCTL, a system call, a subroutine, a function call, an API routine, a remote procedure call, etc., upon initialization of each packet engine, or can be defined by the aggregator **610** upon initialization of the aggregator. The memory allocations can reside in kernel space or in user space. In some embodiments, memory allocations are defined

57

in kernel space and mapped to user space, or vice versa. The memory allocations can be local to the appliance 200, e.g., in cache or main memory of the appliance. In some implementations, memory allocations are not local to the appliance, e.g., reserved on a storage device in communication with the appliance 200 over a network 104.

In various embodiments, the memory allocations are shared between the packet engines and the aggregator 610, e.g., both the aggregator and a packet engine 548 C can write and/or read data to and from a shared allocation 122D. In some embodiments, each packet engine has its own memory allocation, e.g., engine 548A accesses only allocation 122B. As such, no two packet engines share a memory allocation, as depicted in FIG. 6A. In some embodiments, two or more packet engines share a memory allocation. Sharing of memory allocations by packet engines can reduce the amount of memory reserved for performance and trace data.

In various embodiments, the performance and trace aggregating system 600 further comprises an aggregator 610. The aggregator can comprise firmware, software or any type and form of executable instructions executing on a core of the multi-core appliance 200. In some embodiments, the aggregator 610 comprises executable instructions in distributed operation on plural cores of the appliance 200. When configured in distributed operation, communication between the cores by core-to-core messaging, socket connections, or any other mode of data interface can be used to coordinate operation of the aggregator 610. In some embodiments, the aggregator 610 comprises executable instructions executing on a processor dedicated for the aggregator. The dedicated processor can be local to the appliance 200 or located apart from the appliance and in communication with the appliance over a network 104.

In some embodiments, the system 600 comprises two aggregators (not shown). One aggregator, e.g., a performance aggregator, can consolidate performance data obtained from a plurality of packet engines 548A-B to produce unified performance data 606A. One aggregator, e.g., a trace aggregator, can consolidate trace data obtained from a plurality of packet engines 548A-B to produce unified trace data 606B. In certain embodiments, the performance and trace aggregators operate independently. For example, the aggregators may not communicate with each other. In certain implementations, the performance and trace aggregators are in communication with each other. In some embodiments, the performance and trace aggregators are configured to operate on a same core of the multi-core appliance 200. In some embodiments, the performance aggregator operates on a first core, and the trace aggregator operates on a second core of the multi-core appliance 200.

For embodiments where the appliance 200 has a performance aggregator and a separate trace aggregator, two sets of memory allocations 122B-N can be established for each packet engine 548A-B. A first set of memory allocations can store performance data from each packet engine, and the second set of memory allocations can store trace data from each packet engine. In such an embodiment, the performance aggregator shares access to the first set of memory allocations with the packet engines, and the trace aggregator shares access to the second set of memory allocations.

In certain embodiments, only one set of memory allocations 122B-N is established for an appliance 200 having a performance aggregator and a trace aggregator. In such an embodiment, the performance aggregator, trace aggregator, and packet engines can communicate with each other through the shared memory allocations. Communications can be established by the setting of data flags, with or without

58

accompanying commands, in the shared memory allocation by each device. For example, the performance aggregator can set a first flag to indicate a request for performance data. Each packet engine can be configured to recognize the first flag and provide the requested data to its memory allocation, and set a second flag indicating the data is available. The trace aggregator can be configured to recognize the first flag and second flag and not set any flags while the request from the performance aggregator is pending. The performance aggregator can read data from the memory allocations and set a third flag indicating that its request has been completed. The trace aggregator can recognize the third flag and set a fourth flag in the shared memory indicating a request for trace data. Each packet engine can recognize the fourth flag, provide the requested data, and write a fifth flag to the shared memory indicating the data has been written. The performance aggregator can recognize the fourth and fifth flags and not set any flags while the request from the trace aggregator is pending. The trace aggregator can read data from the memory allocations and set a sixth flag indicating that its request has been completed. When the third or sixth flags are present in the shared memory allocations 122B-N, either the performance aggregator or trace aggregator may set a flag in the shared memory requesting data. In this manner, the memory allocations 122B-N can be shared by the performance aggregator and trace aggregator, as well as by the packet engines 548A-N.

In various embodiments, the aggregator further comprises a buffer 122A for storing data obtained from the packet engines 548A-N via the shared memory allocations 122B-N. The data obtained from the packet engines can be stored directly into the buffer 122A, or processed before storing into the buffer. Data stored in the buffer 122A can be exported, or processed and exported to agents 620 requesting data from the aggregator 610. In certain embodiments, buffer 122A comprises a circular buffer, e.g., a circular page buffer, of sufficient size to store a selected amount of operational data or web content obtained from each packet engine. In some cases, the size of the aggregator's buffer 122A is substantially equivalent to the total size of all shared memory allocations 122B-N. In some embodiments, the buffer 122A is mapped by the aggregator upon its initialization, e.g., using IOCTL's, system calls, subroutines, function calls, API routines, remote procedure calls, etc. In some embodiments, the buffer 122A is divided into two sections, one for performance data and one for trace data. In such embodiments, additional space can be reserved for trace data.

In various embodiments, the aggregator 610 produces unified performance data 606A or unified trace data 606B (hereafter sometimes generally referred to as "unified data" 606). The unified data 606 can be exported to the interface agent 615 for forwarding to one or more agents 620 requesting performance or trace data from the appliance 200. In various embodiments, the unified data does not identify the plurality of cores 505A-N and packet engines 548A-N of the appliance. As an example, an agent 620B receiving unified data 606 is unaware of the number of cores 505 or packet processing engines 548 in operation on the appliance 200. The unified data 606 can comprise a data structure having multiple records and multiple data types. The unified data 606 can be stored in any data format on any memory device accessible to the aggregator. In some embodiments, the unified data 606 is not stored by the aggregator, but exported to one or more agents 620. The unified data 606 can include performance statistics and trace data, described above, but representative of the appliance 200 as though the appliance comprised a single core. As an example, unified data 606 can include any

aspect and type of data about the appliance's operation as a virtual server (VIP 275), data about servers to which the appliance is currently maintaining active connections or has had connections to, data about a service managed by the appliance, and/or trace data as described above in relation to the packet engines. In various embodiments, the unified performance or trace data 606 is representative of the appliance's network performance or web content handled by the appliance as a whole. In some embodiments, the unified performance data is a collective representation of the local performance data from each of the packet engines and/or cores. In some embodiments, the unified performance data is a collective representation of the local trace data from each of the packet engines and/or cores.

By way of an example, the aggregator 610 can produce the following unified performance data 606A for a multi-core, multi-packet engine appliance 200:

```
VIP(141.128.58.149:80:UP:WEIGHTEDRR): Hits
(38200495, 18/sec) Mbps(1.02) Pers(OFF) Err(0)
Pkt(186/sec, 610 bytes) actSvc(2) DefPol(NONE)
Conn: Clt(253, 1/sec, OE[252]) Svr(1)
SER1(141.128.49.40:80:UP) Hits(9443063, 4/sec,
P[2602342, 0/sec]) ATr(5) Mbps(0.23) BW1mt(0 kbits)
RspTime(112.58 ms)
Other: Pkt(36/sec, 712 bytes) Wt(10000) RHits(31555)
Conn: CSvr(42, 0/sec) MCSvr(20) OE(16) RP(11) SQ(0)
SER2(141.128.49.39:80:UP) Hits(9731048, 4/sec,
P[2929279, 0/sec]) ATr(9) Mbps(0.27) BW1mt(0 kbits)
RspTime(161.69 ms)
Other: Pkt(41/sec, 756 bytes) Wt(10000) RHits(31555)
Conn: CSvr(32, 0/sec) MCSvr(19) OE(13) RP(4) SQ(0)
```

In this example, the unified performance data comprises a record about virtual service "VIP" provided by the appliance 200 and two records about two servers "SER" providing services at the time the performance data was requested. In the example, the VIP record contains the following information or statistics: (IP address:port:state of the VIP:load balancing method configured for the VIP); (number of hits handled by the VIP in a time interval, current rate of hits on the VIP); (rate of flow of traffic managed by the VIP in Mbps); (state of persistence configured for the VIP); (number of times an error page was generated by the VIP in the time interval); (number of packets per second passing through the VIP, average size of the packets flowing through the VIP); (number of active services that are bound to the VIP); (an indication of whether a default load balancing method is active); (the number of current client connections to the VIP, the current rate at which client connections are established or terminated); (the number of server connections from the VIP in an established state); and (the number of servers actively providing service via the VIP). In the example, each server record contains the following information or statistics: (IP address:port:state of the service); (number of hits directed to the service in a time interval, current rate of hits, [number of hits directed to the service during the time interval due to configured server persistence, current rate of hits due to configured server persistence]), (number of active connections to the service); (rate of traffic flow passed to the service in Mbps); (a bandwidth limit defined by the service); (an average response time of the service); (an average rate of traffic flow in terms of packets during the time interval, average size of the packets); (a weight index used in a load balancing algorithm); (a running hits counter used in round-robin load balancing); (the number of connections to the service, the current rate at which connections are established or terminated); (a maximum number of connections to the service that has been reached); (a number of connections to the service in

an established state); (a number of connections to the service residing in a reuse pool); (a number of connections to the service waiting in a serge queue). In this example, each of the values reported in a record can be obtained by totaling, averaging, or taking a weighted average of values from each packet engine 548A-N. For example, the number of client connections (253) to the VIP can be determined by the aggregator 610 by summing client connections obtained from each of the packet engine's performance record 605PA-N. In some embodiments, multiple occurrences of an aspect are singly counted by the aggregator 610 in the reported value in the unified data. As such, a particular client having multiple connections with the VIP via different packet engines may only be counted as one client connection. Furthering the example, an average response time of a service can be determined by the aggregator 610 by taking a weighted average of response times obtained from each packet engine's performance record 605PA-N. As such, each packet engine's recorded response time for the service can be weighted by the number of hits directed by each packet engine to the service.

The aggregator 610 can also accumulate and provide unified trace data 606B. In various embodiments, each packet engine 548A-N provides its own record of trace data 605TA-N, which can be written to shared memory allocations 122B-N. The aggregator 610 can read each packet engine's trace data and compile an aggregated or unified trace data structure 606B. In some embodiments, the unified trace data 606B comprises a concatenation of trace data from each packet processing engine 548A-N. The concatenation can be compiled according to one of several methods. The method used for concatenation can be specified during system development time or specified in an agent request for unified trace data.

In certain embodiments, unified trace data 606B comprises trace data from each packet engine concatenated by the core for the entire time interval of the trace. By way of an example in which a one-second time interval was specified for the trace and three packet engines 548A-C were actively passing web content during the interval, the unified trace data comprises the web content passed by a first packet engine 548A, followed by the web content passed by a second packet engine 548B, followed by the web content passed by a third packet engine 548C, during the one-second interval. Although the concatenated trace data might represent a three-second time frame, the unified trace data can be reported in "compressed" time, e.g., time scaled by the number of packet engines passing data at the same time.

In certain embodiments, unified trace data 606B comprises trace data from each packet engine concatenated by core for sub-intervals of the entire time interval of the trace. In the above example, the concatenation method can be employed at each 100 ms, 10 ms, 1 ms, or any other sub-interval of time. As such, the trace data recorded by each packet engine would be interleaved when compiling the unified trace data 606B.

In some embodiments, unified trace data 606B comprises trace data substantially assembled according to streams of data exchanged between a particular server and a particular client. In such embodiments, the aggregator 610 can group and concatenate packets of web content according to source and destination IP addresses.

In certain aspects, the aggregator 610 communicates with each packet engine 548A-N using the shared memory allocations 122B-N. Communication can be carried out by setting and detecting the status of data flags in the shared memory. As an example, the aggregator 610 can communicate a request for performance data or trace data from a particular packet engine 548C by setting a data flag in the shared memory

61

122D. The packet engine 548C can monitor the memory 122D for the data flag, and upon detecting the flag write to the shared memory the requested data. After completing the writing, the packet engine can change the data flag or write a new data flag to the memory in a different location indicating the requested data has been written to the shared memory. The aggregator 610 can also monitor the memory 122D for a change in status of the flag, or the addition of a new flag. Upon detecting a changed or new flag, the aggregator can read the data from the shared memory 122D.

It will be appreciated that a flag set by the aggregator can include information identifying the requested data, e.g., performance data, certain elements of performance data, trace data, certain filtered elements of trace data. In some embodiments, the flag itself identifies the requested data or command to be carried out by each packet engine 548A-N. For example, a first flag can identify a request for all performance data, a second flag can identify a request for a first subset of elements of performance data, e.g., elements of virtual service, a third flag can identify a request for a second subset of elements of performance data, e.g., elements of a first service or a first server, a fourth flag can identify a request for trace data, etc. In such embodiments, each packet engine 548A-N can be configured to recognize the one or more flags and identify the requested elements from the flag itself. In some embodiments, the flag is accompanied by data written in the shared memory by the aggregator 610, wherein the accompanying data identifies the requested data or command to be carried out by each packet engine. It will be appreciated that the aggregator 610 can set the same flags in all memory allocations 122B-N shared with the packet engines 548A-N, or can set different flags in the memory allocations.

When shared memory is further shared by packet engines, e.g., packet engines 548A and 548B share allocation 122B, communication via flags and time multiplexing can be employed to manage data written to and read from the shared memory. As an example, the aggregator 610 can request sequentially data from the two packet engines 548A and 548B in this example. A first flag identifying a first packet engine 548A and a data request can be written by the aggregator to the shared memory 122B. Packet engine 548A monitoring the memory can detect the flag, write the requested data to the memory 122B, and alter the flag or add a new flag to indicate it has written the data to the memory, while packet engine 548B ignores the flag since the flag does not identify engine 548B. Aggregator 610 can read the data provided by the first packet engine 548A from the memory 122B, and write a second flag identifying the second packet engine 548B and data requested. The second flag can replace the first flag or altered flag. Packet engine 548B can detect the flag, write the requested data to the memory 122B, and alter the flag or add a new flag to indicate it has written the data to the memory, while packet engine 548A ignores the flag. Aggregator 610 can then read the data provided by the second packet engine 548B.

In various embodiments, aggregator 610 is adapted to communicate with agents 620. The communication with agents 620 can be via an agent interface 615. The agent interface can comprise software, firmware or any type and form of executable instructions executing on a core of the appliance. In some embodiments, the agent interface 615 comprises a set of instructions which are part of the aggregator 610. In some embodiments, the agent interface 615 comprises a set of instructions operating separately from the aggregator 610. In some instances, the agent interface is in operation on the same core as that for the aggregator, and in some cases the agent interface 615 is in operation on a different core from that on

62

which the aggregator 610 operates. In various embodiments, the agent interface 615 communicates with the aggregator 610 via socket connections.

In certain embodiments, the agent interface 615 comprises two IP ports, e.g., port 3333 and port 5555, and software associated with establishing and maintaining socket connections on these ports. One port can be used for receiving requests from and transmitting data to one or more agents 620, such as the data relating to performance statistics. One port can be used for receiving requests from and transmitting data to one or more agents, such as the data relating to trace data. In certain embodiments, the aggregator 610 is configured to listen for socket connections on ports of the agent interface 615. In some embodiments, the agent interface comprises one or more data files. As an example, data requests can be written to a "data request" file by an agent 620A and read from the file by the aggregator 610. The aggregator can write the requested data to a "data return" file, which can be read by the agent. In some embodiments, the agent interface 615 comprises software which translates requests received from one or more applications into a format recognizable by the aggregator 610.

An agent 620A, 620B (generally 620) can comprise any entity requesting performance or trace data from the appliance 200. An agent can be a system operator or administrator entering commands in a command line interface (CLI) in communication with the appliance. An agent can comprise a graphical user interface, dashboard, proprietary software, or publicly available software, e.g., TCPDUMP, Wireshark, NETSVC, SNMP, etc., configured to communicate with the appliance 200. In some embodiments, an agent comprises another networked device, e.g., a server 106, a client 102, or another appliance 200', which periodically interrogates the appliance 200 for performance or trace data. In certain embodiments, an agent is a connection, module or agent of a client for interfacing with and/or gathering information from the aggregator.

Referring now to FIG. 6B, an embodiment of method 650 for use with the performance and trace aggregating system 600 is depicted. In various embodiments, the method 650 comprises collecting or capturing 652 data by each packet engine, storing 664 data to a buffer by the aggregator, consolidating 666 the data by the aggregator, receiving 668 a request for data by the aggregator, and transmitting 670 the consolidated data by the aggregator in response to the request. This and other steps-based depictions in this disclosure are not intended to portray a specific sequence of the steps, and certain steps can be performed before or after other steps. The method can comprise a subset of the steps depicted, and in some embodiments include additional steps as described below.

In certain embodiments, the step of collecting 652 comprises collecting by each packet engine performance data identifying statistics of a service provided by a server for which each packet engine manages network traffic. In some embodiments, the performance data identifies statistics of a server or service only associated with a current or active connection between each packet engine and the server providing the service. In some embodiments, the performance data identifies statistics of a server or service associated with a past connection to a server. In certain embodiments, the collected performance data identifies statistics of virtual service provided by each packet engine. The performance data can include any element of information about virtual service of each packet engine as described above, e.g., the IP address and port configuration of the VIP, the state of the VIP, etc. The performance data can include, alternatively or in addition to

63

the VIP data, any element of information about the server or service of the server as describe above, e.g., the IP address and port configuration of the server, the state of the service provided by the server, etc. In some embodiments, the step of capturing **652** comprises capturing, by each packet engine, trace information for network traffic received or transmitted by each packet engine. Trace data can include, without being limited to, any element from the following list: requests for web content, responses to requests for web content, encrypted web content, non-encrypted web content, and any combination thereof as described above. In some embodiments, the collecting **652** of performance data is done periodically, e.g., every 10 seconds, every 7 seconds, every 3 seconds, every 1 second, by each packet engine. In various embodiments, the step of collecting **652** further comprises determining, by each packet engine **548A-N**, performance statistics and recording the statistics to performance records **605PA-N**. In various embodiments, the step of capturing **652** comprises recording, by each packet engine, trace data to trace records **605TA-N**. In some embodiments, collecting **652** of performance data is done in response to a request for performance data indicated by the aggregator **610**. In certain embodiments, the capturing **652** of trace data is done in response to a request for trace data indicated by the aggregator **610**.

Further aspects of an embodiment of the step of collecting or capturing **652** are depicted in FIG. **6C**. In certain embodiments, the step of collecting or capturing **652** comprises sending or storing **654a** command and setting **655** a flag by the aggregator in each shared memory allocation **122B-N** for each of the plurality of packet engines. In some embodiments, the data flag comprises an indicator and a command. The indicator can identify a status of the request, e.g., data requested, no data requested, data requested at a specified time, data requested for a specified time interval, data requested for a particular packet engine, etc. The command can comprise details about the elements of data requested as well as specific types of actions, e.g., provide virtual service data, dump all data immediately, terminate data collection, terminate data capturing, begin data capturing, filter data in accordance with specified filtering criteria. In various embodiments, the step of collecting or capturing **652** further comprises changing, by each packet engine, the data flag in shared memory allocations **122B-N** when each packet engine completes the step of collecting or capturing **652**.

The step of collecting or capturing **652** can further comprise detecting, by each packet engine, the data flag set by the aggregator **610** in memory shared between the packet engine and the aggregator. Each packet engine **548A-N** can monitor the shared memory for the data flag. Each packet engine can further interpret any data, e.g., a command set by the aggregator **610**, included with the flag and identifying certain aspects of data requested, e.g., performance data, certain elements of performance data, trace data, filtered elements of trace data, time intervals associated with the data. Data included with the flag can also include IOCTLS, objects, data structures, function calls, procedure calls, routines, macros, and the like.

The step of collecting **652** can comprise executing, by each packet engine, diagnostic routines to compile performance statistics for each packet engine. The diagnostic routines can identify performance data or certain elements of performance data as requested by the aggregator **610**. In some embodiments, the routines can comprise maintaining counters which determine any of the following types of values, total number, maximum number, minimum number, average number, frequency, etc. The step of collecting **652** can comprise collecting data for current active connections between active clients

64

and servers actively providing service. In some embodiments, the step of collecting **652** can comprise collecting data historical data of performance statistics stored in an archive accessible to the packet engine. The step of capturing **652** can further comprise executing routines by the packet engine to record during a selected time interval network traffic passed by the packet engine, e.g., record trace data during a specified time interval. In some embodiments, the capturing further comprises filtering the captured data. The filtering can comprise recording, by each packet engine, only network traffic which meets a filter criterion, e.g., from a particular source IP address, during a particular time interval, to a particular destination IP address, etc.

In various embodiments, the step of collecting or capturing **652** further comprises collecting or capturing data by each packet engine for a selected time interval. The selected time interval can be identified by the aggregator **610**, and established with the data flag and/or command set by the aggregator in shared memory **122B-N**. In certain embodiments, a default time interval is used for instances when no time interval is specified by the aggregator. In various embodiments, the step of collecting or capturing **652** further comprises writing, by each packet engine, performance data or trace data to the shared memory allocations **122B-N**. The performance data can be written, by each packet engine, to the shared memory allocations **122B-N** from each packet engine's performance record **605PA-N**. The trace data can be written, by each packet engine, to the shared memory allocations **122B-N** from each packet engine's trace record **605TA-N**.

Referring again to FIG. **6C**, the step of collecting or capturing **652** can further comprise monitoring, by the aggregator, the status of the data flags set in each shared memory allocation **122B-N** to determine **656A** whether each flag's status has changed. If it is determined that a flag's status has changed, the aggregator **610** can read **657** performance data or trace data from the shared memory allocation indicated by the changed flag status. If it is determined that a flag's status has not changed, the aggregator **610** can determine **656B** whether a packet engine is slow in providing the requested data, e.g., the packet engine has not provided the requested data within a specified time window. If it is determined that the packet engine is not slow, the aggregator can execute **658** another task and return to checking **656A** the flag status at a later time. If it is determined **656B** that the packet engine is slow in providing the data, the aggregator **610** can issue a command which forces **659** the packet engine to dump its performance data or trace data. The aggregator **610** can then read **657** the dumped data. In cases where a forced dump is executed by a packet engine and the packet engine has not complete all of its data acquisition, any remaining data values to be acquired or any remaining data trace to be acquired can be represented as zero values when the packet engine writes the acquired data to the shared memory allocation.

The step of storing data **664** can comprise storing to a buffer **122A** by the aggregator **610** performance data or trace data from each packet engine. In some embodiments, the step of storing **664** comprises storing, by the aggregator, performance or trace data to a file or location external to memory **122**. In various embodiments, the step of storing can include detecting, by the aggregator, a change to a data flag set in shared memory allocations **122B-N** by the aggregator. If the aggregator detects a change in a data flag for a particular memory allocation **122C**, the aggregator can read data from the shared memory having a changed flag and write the data to the aggregator's buffer **122A**. If the aggregator does not detect a change in a data flag for a particular memory allocation **122B**, the aggregator can attend to other processing tasks,

65

e.g., detect the status of a data flag for the next shared memory allocation, and return to check **656A** the data flag for the particular memory allocation **122B** at a later time.

The step of storing **664** can comprise writing data from a shared memory allocation directly to the aggregator's buffer **122A**. In some embodiments, the step of storing **664** further comprises processing the data from the shared memory allocation before writing the processed data to the aggregator's buffer. In some embodiments, the step of storing **664** comprises reading data from shared memory allocations **122B-N**, processing the data, and exporting results from the processed data directly to an agent without storing the data in the aggregator's buffer **122A**. In some embodiments, the processing of data comprises consolidating the data obtained from the plurality of packet engines.

In some embodiments, the step of storing **664** further comprises maintaining data structures representative of per core statistics and data structures representative of global statistics. For example, the aggregator **610** can maintain a per core data structure which identifies individual services currently managed by each core. The aggregator **610** can further maintain a global data structure which identifies on how many cores a particular service is handled. In certain embodiments, when a new service is added on any packet engine, the packet engine dumps its performance data to its shared memory allocation. The aggregator **610** can then detect the addition of a service and update its per core and global data structures. In some embodiments, the addition of a new service is detected when each packet engine writes its performance or trace data to shared memory in response to a request for performance or trace data. In some implementations, when a service is deleted from any packet engine, the packet engine dumps its performance data to its shared memory allocation. The aggregator **610** can then detect the deletion of a service and update its per core and global data structures. In some embodiments, the aggregator **610** removes the service from all per core data structures and the global structure. In some embodiments, the aggregator **610** only removes the service from the global data structure when all packet engines indicate that they no longer manage the service.

In various embodiments, methods **650** for aggregating performance or trace data comprise consolidating **666**, by the aggregator, the performance or trace data to provide unified performance data **606A** or unified trace data **606B** of the multi-core, multi-packet-engine appliance **200**. The unified data can be representative of the appliance as though the appliance comprised a single core. As an example of consolidation, a case where two packet engines **548A**, **548B** provide service to a plurality of clients **102A-N** from the same server **106A** is considered. A first packet engine **548B** may record in shared memory **122C** that it has 3 active connections with server **106A** for the service and that its traffic volume to the service is 85 packets per second. The second packet engine **548A** may record in shared memory **122B** that it has 4 active connections with server **106A** for the service and its traffic volume to the service is 37 packets per second. For this example, the unified data **606** obtained from the step of consolidating **666** would represent the appliance **200** as having 7 active connections with server **106A** for the service and a traffic volume of 122 packets per second.

In some embodiments, the step of consolidating **666** comprises processing performance data **605P** or trace data **605T** obtained from the plurality of packet engines **548A-N**. The processing can include executing macros, functions, routines, or sets of instructions to compile unified performance statistics **606A** or unified trace data **606B** representative of the appliance **200**. As an example, the processing can comprise

66

totaling a number of connections to distinct servers, totaling traffic volume to each distinct server, totaling a number of hits on each distinct server, obtaining averages or weighted averages of traffic volume to each distinct server, interleaving trace data to provide a unified trace, etc. as described above in reference to unified performance data **606A** and unified trace data **606B**. In various embodiments, the aggregator **610** consolidates the data from the plurality of packet engines **548A-N** after the data is written to buffer **122A**. The step of consolidating can further comprise writing the consolidated data **606** to a separate location in memory **122** or to a location reserved in the buffer **122A** for consolidated data. In some embodiments, the aggregator consolidates **666** the data from the plurality of packet engines after reading the data from shared memory allocations **122B-N**, and writes the consolidated data **606** to buffer **122A**, or exports the unified data **606** directly to an agent **620A** or to a file external to memory **122**. In certain embodiments, the step of consolidating **666** is repeated periodically, e.g., at time intervals between about 50 milliseconds (ms) and about 500 ms, between about 500 ms and about 5 seconds, between about 5 seconds and about 60 seconds.

When consolidating **666** trace data, the aggregator **610** can further represent the consolidated data in terms of actual processing time, compressed time, or both. As an example, the processing by the multi-core appliance of a stream of packets from a particular server is considered. In various embodiments, the stream of packets flows through and is processed by the multi-core appliance **200**. In some embodiments, the packets are parsed to different packet engines which can process their allotted portions of the stream at the same time or different times from the other packet engines. A trace representation of the packet stream in terms of actual processing time would comprise consolidating the trace data to reflect the total amount of processing time of each core used in processing the stream of packets, e.g., concatenating each core's processing segment. A trace representation of the packet stream in terms of compressed time would comprise consolidating the trace data to reflect the time at which each portion of the stream was processed by the appliance. Compressed time could provide an indication of multi-packet-engine processing, e.g., a high rate of content handling for a particular time interval (multi-processing) and a low rate of content handling at another time interval (single-core-processing) for a stream of packets.

In some embodiments, the step of consolidating **666** further comprises consolidating the number of connections to a service from the plurality of packet engines **548A-N** and providing a total number of connections to the service from the appliance **200**. The service can be provided by one or more servers **102** over the network **104**. The total number of connections can be a raw total or represent a number of connections from distinct clients. (Some clients may have more than one connection to a service.) In some embodiments, the consolidating **666** further comprises consolidating the average response to a service from the plurality of packet engines **548A-N**, and providing a unified average response time to the service from the appliance **200**. The unified average response time can be calculated in any one of several ways, e.g., averaging average response times obtained from each packet engine, calculating a weighted average such as weighting the average response obtained from each packet engine by the number of connections made to the service by the packet engine.

In some embodiments, the step of consolidating **666** further comprises consolidating the number of bytes passed for a service from the plurality of packet engines **548A-N**, and

67

providing a total number of bytes passed for the service by the appliance 200. In some implementations, the step of consolidating 666 further comprises consolidating the number of different servers obtained from the plurality of packet engines 548A-N, and providing as a portion of unified data a number of distinct servers providing network service via the appliance 200. The number of different servers may include all active servers on each packet engine, all recently active servers, servers in a reserve or reuse pool, all past servers, or any combination thereof. The number of distinct servers can comprise a subset of the different servers having distinct IP addresses.

The step of receiving a request 668 can comprise receiving, by the aggregator, a request for performance data or trace data of the appliance. The request can be received from one or more agents 620. The request can be received by the aggregator through an IP port, e.g., via an established socket connection managed by the agent interface 615. In certain embodiments, the step of receiving a request 668 comprises listening 653 on one or more IP ports for a socket connection and receiving data representative of a request through an IP port. The step of receiving 668 can further comprise monitoring activities on IP ports managed by the agent interface 615. In certain embodiments, the step of receiving a request 668 comprises monitoring a data file or location in memory for a newly entered request and reading the request from the file or location in memory. In various embodiments, the request comprises data identifying elements of performance statistics or trace data desired by the entity issuing the request, and identifies a destination for transmitting a response to the request. In some implementations, a request comprises command line code, IOCTL's, objects, function calls, or any type and form of executable instructions. The step of receiving 668 a request can comprise executing, by the aggregator, the received instructions, or providing the received instructions, or a portion thereof, to each packet processing engine 548A-N.

In various embodiments, the aggregator 610 is adapted to transmit 670 consolidated or unified data 606 in response to a received request for the data. The unified data 606 can be transmitted in any format. In particular embodiments, the step of transmitting 670 further comprises converting the unified data 606 to a format specified by the agent 620 issuing the request. In certain embodiments, the step of transmitting 670 comprises exporting the unified data 606 from buffer 120A to agent interface 615, which in turn can forward the unified data to the agent 620 requesting the data. In some embodiments, the step of transmitting comprises identifying, by the aggregator, a memory location holding the unified data 606. The memory location can be identified to the agent interface and/or the agent requesting the data whereupon either the agent interface or agent requesting the data can read the unified data 606 from the identified memory location.

G. Systems and Methods For Scalable N-Core Stats Aggregation

Referring now to FIG. 6D, one embodiment of a system for aggregating data or statistics from cores of a multi-core appliance is depicted. In brief overview, the system includes a plurality of cores 548A-N, shared memory 112, and an aggregator 610. In various embodiments, one or more clients may make a request to the aggregator for data or statistics (e.g., performance statistics and trace data) associated with one or more cores of the multi-core system. The depicted system may be implemented as a variation, combination, reconfiguration and/or improvement over any of the embodiments of features described above in connection with FIG. 6A.

68

In certain embodiments, the agent interface and/or agent allows a client to access data or statistics (hereafter sometimes generally referred to as "data" or "statistics") collected by the aggregator 610. A client may be any network device (e.g., client 102, server 106, appliance 200, etc.), service, virtual machine, virtual service or program that uses the statistics or provides the statistics to another entity either processed or unprocessed. In certain embodiments, a client may be a device transmitting and/or receiving traffic via one or more cores or packet engines of the appliance. In some embodiments, a client may be an administrative service, monitoring service or other type of service, e.g., an nsnetshvc associated with CITRIX NETSCALER. In other embodiments, nsnetshvc may be an agent interface for the aggregator and an agent of a client. In still another embodiments, nsnetshvc may be an agent of a client. A nsnetshvc, as a client and/or client agent, may generate and/or deliver commands for requesting statistics from the aggregator. A nsnetshvc, as an agent, may deliver or convey client commands for requesting statistics from the aggregator.

In some embodiments, a client may request certain events to be reported in real-time or substantially in real-time. These events may be represented by statistics of packet processing engines, virtual servers and/or associated member services processed and/or tracked by one or more cores of the multi-core system. A client may also request for statistics from performance records, e.g., using a stat command. A client may include or provide a dashboard or other interface for organizing, summarizing and/or presenting statistics to a user or program. In some embodiments, a dashboard may process and/or provide data in real-time or substantially in real-time. A client request may comprise a request for statistics associated with two or more cores of a multi-core system, or statistics for the multi-core system as a single unit. In one embodiment, for example in the NETSCALER context, a client may use a nsconmsg and/or nscollect command in a request. Another client request may comprise a request for statistics associated with a particular core of a multi-core system.

The system may support and interface with any number of different types of clients. In some embodiments, and by way of illustration, the different types of clients may be categorized into at least (1) streaming or non-filtered clients, and (2) filtered request-response clients. Streaming or non-filtered clients may transmit requests that are sometimes referred to as kernel mode packet engine (KMPE) type requests. Such requests may include commands such as nsconmsg. These requests may communicate with an aggregator on a particular port (such as 3333), instead of sending IOCTL calls to the kernel of the appliance. The aggregator may be assigned to and/or listening on the particular port. In some embodiments, streaming clients may parse raw and/or full performance data streams (sometimes generally referred to as "perf-streams" or "data streams") into statistics and/or records. Streaming clients may create and/or maintain the data structures for the statistics and/or records. Thus, streaming clients may have the same or substantially the same access to statistics and information that an aggregator or a packet engine may have. By way of illustration, requests from streaming clients can include nsconmsg, dashboard and nscollect requests.

The aggregator may support one or more filtered request-response clients. Request-response clients may send a request for statistics for one or more particular devices (i.e., vservers, packet engines and/or cores). These devices may already be known or identified by the client. The client may already be in communication with these devices. For example, the multi-core system may have assigned a particular core to process network traffic associated with the client, or the client may

have previously communicated a packet via a specific packet engine or core. Responsive to such a request, the aggregator may filter and/or process the statistics gathered at the aggregator, generate a response and/or send the response responsive to a client request. In some embodiments, however, per- entity, e.g., per-core or per-packet-engine statistics may not be available. For example, per-core statistics may not be available because statistics for two or more cores or packet engines may have already been merged together. By way of illustration, and not intended to be limiting in any way, requests from request-response clients may include stat, snmpd, nsld and show commands or requests.

As discussed above, the aggregator may collect statistics from entities such as packet engines and consolidate the statistics, for example, based on counter definitions. The counter definitions may specify an interval, schedule and/or sequence for collecting statistics from each packet engine. The statistics may be of any type or form, such as summation, average, minimum and maximum per-packet-engine (sometimes referred to generally as “per-PE”) values. The aggregator may create per-PE records based on the collected statistics and may consolidate these records to an internal buffer of the aggregator. The aggregator may parse statistics from each of the packet engines. The aggregator may store or maintain the per-PE statistics prior to consolidation. In certain embodiments, the aggregator consolidates and stores the packet engine statistics into the internal buffer. In some embodiments, one or more clients may access information from the internal buffer. In other embodiments, the aggregator generates a response to a client request using information from the internal buffer. The internal buffer may be or incorporate any type or form of memory device such as cache **140**, storage **128**, memory **122**, **264**, buffer **243** and disks **428** described above in connection with FIGS. **1E-F**, **2A** and **4A**. The aggregator may also perform multi-tasking or interleaving between any of the above tasks to check for client requests and/or respond to client requests.

In some embodiments, the aggregator may allocate memory and/or data structures to maintain or manage mappings between local device identifiers and global device identifiers. Local device identifiers may be assigned on a per-PE or per core basis. In certain embodiments, the aggregator may assign a local device identifier to an entity such as a packet engine, virtual service, virtual machine, service, application or some other packet processing module in a core. In some embodiments, each entity may be assigned one or more device identifiers. For example, each packet engine or core may execute, provide or be associated with one or more devices (e.g., virtual servers executing in the core). For example, a vserver v1 executing in connection with PE-0 may have a local device number **100**, while a similar vserver executing in connection with PE-1 may have local device number that is the same (e.g., 100) or different (e.g., 200) from the vserver v1. Thus, across a plurality of cores, local device numbers of such entities may not be unique. In some of these embodiments, the aggregator translates per-PE or local device identifiers to a global device identifier when consolidating statistics from each packet engine. A global device identifier may uniquely identify an entity across a plurality of cores.

In some embodiments, each core may operate with the aggregator via shared memory **112**. For example, the aggregator may request pages or information from a core by writing the request into the shared memory. The core may receive the

request by accessing the shared memory **112**. The aggregator **610** may read the requested pages or information from the shared memory. The aggregator may process the requested pages or information and store them in an internal buffer of the aggregator, which may be local memory (e.g., of the aggregator, not shared memory). In response to a client request, the aggregator may establish a socket connection with the corresponding client or client agent. The socket connection may be used to read and/or write IOCTL messages, e.g., from the client to a nsnetvc, and from the nsnetvc to the aggregator.

One core may provide data to the aggregator using one or more dedicated memory locations. In certain embodiments, an aggregator may share or reuse one or more memory locations between different cores, e.g., at different times.

The aggregator, in processing the statistics, may have certain memory requirements. For example and in some embodiments, processing absolute records from a packet engine may involve the following:

Collecting stats=100 ms.

Parse the stats=1200 ms.

Dump the stats into client buffer=400 ms.

Processing absolute statistics, which may exclude some records such as devicename, devlink, codeinfo records, may involve the following:

Collecting stats=100 ms.

Parse the stats=1045 ms.

Dump the stats into client buffer=400 ms.

Processing differential statistics (e.g., incremental changes or updates to the statistics) may involve the following:

Collecting stats=10 ms.

Parse the stats=200 ms.

Dump the stats into client buffer=400 ms.

The total number of devices established, created, deployed and/or monitored in a configuration may be $200K=2*10^5$. The total number of device statistics records maintained in the aggregator may be $2M=2*10^6$. The per device counter may be set to 10. Each batch of device statistics may take up at least 24 bytes (current, previous and earlier batches) for example. Earlier batches of data may be used for maintaining rate calculations. The total per-PE statistics space requirement in the aggregator may be $24*2*10^6$ bytes=48 MB. In some embodiments, under normal traffic patterns, per device counters may be assumed to be approximately 20-30 (e.g., instead of 10). According to this scenario, the per-PE statistics memory requirement may be around 100-150 MB or more. In some embodiments, memory requirements may increase, for example, when features such as configuration scaling, core scaling and counter scaling, are added.

In some embodiments, the aggregator may be configured to aggregate statistics without allocating memory for maintaining per-PE statistics. Each packet engine or packet processing entity may be assigned at least one global device identifier, e.g., by the appliance or a server, for transferring or dumping per-PE statistics. The aggregator may store data structures that map local device identifiers to global device identifiers in shared memory with each core. In certain embodiments, the aggregator may use a single block of shared memory shared across a plurality of cores. In some embodiments, the aggregator may allocate or assign specific memory locations or segments for shared access with a particular core. By way of illustration and not intended to be limiting in any way, an embodiment of data structures maintained in shared memory for managing device identifiers is as follows:

```

u32bits nspm_cur_global_devno; /* global devnumber for incrementing */
/* structure for handling the device name <-> devnumber mapping*/
typedef struct nspm_devnamerec
{
    NS_LIST_ENTRY(nspm_devnamerec) dn_list; /* next in chain */
    struct nspm_devnamerec *dn_next_free; /* to maintain the number of devices in the
free list*/
    u32bits dn_devno; /* global device number */
    u08bits dn_devname_len; /* length of this device name */
    u08bits reserved_1[3];
    u08bits *dn_devname; /*Max 192 characters, for example*/
} nspm_devnamerec_t;
typedef struct {
    struct type *lh_first; /* first element */
    u32bits devname_slck; /*lock for this bucket */
} nspm_devnamehead_t;
nspm_devnamehead_t nspm_devnamehashheads[NSPM_DEVNAMEHASHSIZE];
u32bits nspm_getdevnamedevno(u08bits *name);

```

These structures or application processing interfaces (APIs) may be used by a packet engine to obtain a unique mapping between a device (e.g., device name or local device identifier) and a global device identifier. In some embodiments, by using global device identifiers, the aggregator may directly add per-PE differential statistics and/or assign absolute statistics to aggregated values. This may enable the aggregator to bypass certain steps, such as the steps of updating per-PE data structures before updating aggregated values.

In certain embodiments, the aggregator does not maintain per-PE statistics. The aggregator instead dumps the per-PE data streams unchanged to a log file or memory (e.g., newslog or buffer) using the global device identifiers numbers or associated markers to distinguish the data streams, PEs or cores.

In certain embodiments, the aggregator may provide a computing or process thread (hereafter sometimes generally referred to as “thread”) to a packet engine, e.g., for collecting and updating per-PE statistics. A thread may be any type or form of processing unit, task, code, application, agent, module or computing step(s). A thread may execute independently of another thread, or may communicate or interoperate with another thread. A given thread may collect a data or statistics stream from a respective packet engine, e.g., directly or via shared memory. The thread may parse, analyze or process the data stream and may update some portion of aggregated statistics (e.g., of a given device, packet engine or core). The thread may process or track a respective data stream according to a global device identifier. The thread may write the data stream or statistics into a logfile or memory unchanged (e.g., as generated or recorded by a packet processing engine or a monitor executing on a corresponding core) or substantially unchanged. The thread may write the data stream to a logfile or memory without consolidating data from one or more other entities, devices, packet engines or cores. This multi-threading process, supporting multiple entities across one or more cores, may improve overall processing time. For example, processing time may be reduced because the aggregator writes data streams directly into a logfile instead of to an internal buffer requiring proper formats. Concurrent or parallel processing via multi-threading may also reduce processing time.

In some embodiments, the aggregator may receive a request from a client. The aggregator may identify the request as a dashboard request for an aggregated data stream. The aggregator may be built, designed and/or configured to parse, collect and/or consolidate a plurality of data streams. The aggregator may generate an aggregated data stream from one or more data streams or cores. The aggregator may, for

example, generate an aggregated data stream in real time, or in response to receiving a request. Responsive to the request, the aggregator may read or access one or more data streams and/or modify processing on an aggregated data stream.

The aggregator may create and assign a thread to generate, update, handle or process this aggregated data stream. The aggregator may process and/or write at least a portion of the aggregated data stream to the internal buffer. The aggregator, client and/or client agent may establish a dashboard connection to the internal buffer to access the stored data. In some embodiments, the client or client agent may send a message to the aggregator when the dashboard connection is closed or no longer required. Responsive to this message, the aggregator may stop writing or generating the data stream to the internal buffer. Independent of and/or concurrent with this thread, another aggregator thread (e.g., a main aggregator thread) may serve one or more client requests of the same or a different type (e.g., stat, snmpd, nsld, show requests).

In some embodiments, a client may seek access to a data stream handled by the aggregator. The client may initiate the request via a nsconsmg or nscollect request. The corresponding client or agent may access the data stream by connecting to a port of the aggregator (e.g., NSAGG_LISTEN_PORT). The client or agent may issue IOCTL calls via the port (e.g., NSAPI_ADD_RET_PERFDATAHEADER, NSAPI_ADDPERFDATAHEADER, NSAPI_WAITRACEDATA, NSAPI_GETRACEDATA and NSAPI_GETFIRSTTRACE-DATA) to access the internal buffers of the aggregator.

In other embodiments, a client or client agent may access a performance stream from a logfile (e.g., a newslog file) or a storage device that receives or buffers the data stream. The logfile or storage device may reside anywhere on the network, such as on the aggregator or in communication with the aggregator. The data streams or statistics may be retrieved from the logfile or storage device instead of from the aggregator (e.g., from local buffers of the aggregator). The aggregator may directly write at least a portion of the data stream into the logfile or storage device. A network service or agent, for example operating on behalf of the client, may access the data stream from the logfile or storage device.

In some embodiments, a client request may initiate or execute a program or instructions to collect information from one or more data streams. For example, a nsconsmg request may be associated with a set of instructions. These instructions may be compiled into a binary file, executable and/or program file. The compiled instructions may be incorporated or added to a library (e.g., libnsapps). The compiled instructions may reside or be installed in the aggregator, or accessible by the aggregator. Instructions associated with various

types of client requests may be used for debugging purposes, for collecting data streams from the aggregator, and for writing a data stream to a logfile, for example.

In certain embodiments, the instructions may be executed as a process. A process, such as one associated with the nsconmsg instructions, may be made event-aware. Such a process may receive a notification when the logfile or storage device is updated with newer data. The process may parse or access the newer data responsive to the notification. In some embodiments, and by way of illustration, a disk access may occur every 7 seconds for a process to parse new pages added to the logfile or storage device. A large configuration may involve disk accesses of 300 pages per packet engine, for example. On a 7-packet-engine system, for example, 8×300 pages may be accessed, which may translate to about 19 MB per 7 seconds, or about 2.5 MB per second. These are page reads, which may be preferred over socket reads. In certain embodiments, a system having a typical configuration may handle about 1-20 pages per 7 seconds.

Storage requirements may be reduced by not having the logfile or storage device contain an aggregated data stream. In some embodiments, storage space levels or savings of $1/(\text{num of PE}+1)$ with respect to the alternative may be achieved. An aggregated data stream may include data collected in real time and/or from all or some cores and their entities. In some embodiments, portions of data streams (e.g., representative, estimated or sample statistics) are collected from each PE for storage in the logfile or storage device. If aggregated statistics are needed or preferred, the requesting client may obtain the different data streams and perform the aggregation (e.g., at the client or by the client agent). In some embodiments, counter information for doing so can include different types of counters (e.g., `_TNC/_TNAV/_TNMAX/_TNMIN`, etc.). Each counter may add or provide a data construct, format and/or location (e.g., `codevalformat`) for holding or tracking this information. These counters may be implemented in or by the aggregator.

In some embodiments, a counter may be used to partition and/or identify a data stream from one or more PEs. Data streams from each PE may be interleaved using the counter, e.g., according to a predetermined sequence and/or time period for data collection/logging. In certain embodiments, a counter may generate a marker (e.g., a count or identifier) for insertion with a particular data stream. The marker may be used to uniquely identify a data stream from a particular source entity (e.g., core, packet engine or virtual server). In some embodiments, a marker is generated based on (e.g., to uniquely correspond to) a global device identifier of an entity. A marker may, for example, include a global device identifier, or be uniquely mapped or translated into a global device identifier using a mapping table or other means.

In some other embodiments, an aggregated stream may be stored or maintained in the logfile or storage device. This may reduce the amount of time for a client to obtain consolidated values. Moreover, in certain embodiments that store or maintain an interleaved or aggregated stream in the logfile or storage device, markers for the respective aggregated streams may be used for jumping or identifying between data of different streams, such that parsing time within the logfile or storage device can be saved.

In some embodiments, the aggregator may adaptively inform or communicate with a packet engine to adjust the data stream transfer, dumping interval or schedule. For example, there may be situations in which some aggregator threads are operating slower than the data stream dumping interval in the packet engine, such as when the processor handling the multi-threads have higher loads or is handling more threads than

usual. The dumping interval may be decreased or increased to synchronize the packet engine data transfer with the aggregator threads. In some embodiments, the aggregator may use flags or other monitoring means to track operating speed and/or transfer readiness, for example, using the flags described above in connection with FIGS. 6A-C.

The aggregator may adaptively adjust the transfer interval based on the logfile generation rate. In some embodiments, a packet engine may dump 300 pages every 7 seconds. Based on an eleven-PE, for example, logfile generation may involve the following: Absolute records without statistics (X) may cover about 1800 pages, which translates to about $1800 \times 8 \times 1024$ byte or 14 MB of memory. Absolute stats from one PE (Y) may cover about 1200 pages, which translates to about 9.37 MB. Differential stats from one PE (Z) may cover about 300 pages, which translates to about 2.34 MB. Based on the above, one set of absolute record for clients on an 11-PE system may translate to about $X + (11+1)Y$ or 126.44 MB. An extra unit of memory (Y), as indicated, may be allocated to account for the aggregator. In some embodiments, absolute statistics (e.g., dumped when a configuration change happens) may translate to about $(11+1)Y$ or 112.44 MB. Differential statistics (e.g., at every 7 seconds) may translate to about $(11+1)Z$ or 28 MB. In this case, a differential record may consume about 4 MB per second. If, the logfile is compressed (gzipped) when it reaches a threshold, e.g., 40 MB, compression may be triggered every 70 seconds. Thus, the aggregator may adaptively or accordingly adjust the rate of transfer or data stream dump between the packet engines, threads and/or the logfile. For example, if the aggregator determines that the logfile size is incrementing too fast (e.g., beyond a threshold rate), the aggregator may lengthen the interval between transfers/dumps.

Referring now to FIG. 6E, an embodiment of steps of a method 680 for aggregating performance statistics from cores of a multi-core system is depicted. In brief overview, a multi-core system intermediary between one or more clients and servers maintains in shared memory of the multi-core system a global device number for each core of the multi-core system (681). Each core may include one or more packet engines processing network traffic between the one or more clients and servers. The multi-core system may provide, via an aggregator, a computing thread for each core of the multi-core system to gather data from the corresponding core (683). A first computing thread of the aggregator may collect statistics of network traffic processed by one or more packet engines on a first core (685). The first computing thread may transfer the statistics with a marker to a statistics log of the multi-core system (687). The marker may correspond to a global device number of the first core. The multi-core system may adaptively reschedule the transfer by monitoring the operation of each computing thread (689). In further details of step 681, a multi-core system intermediary between one or more clients and servers maintains in shared memory a global device identifier (or number) for each core of the multi-core system. Each core may include one or more packet engines processing network traffic between the one or more clients and servers. An aggregator of the system may provide and/or maintain the shared memory. The aggregator may provide and/or maintain a block of shared memory for each core, or for a plurality of cores. The aggregator may provide and/or maintain a block of shared memory for use with one or more packet engines, virtual machines, applications, services and/or other entities of each core. For example, the aggregator may provide and/or maintain a respective portion of the shared memory for use with each packet engine or other entity of a core.

The aggregator may generate, provide and/or assign one or more global device identifiers to each core or packet engine. The aggregator may generate, provide and/or assign one or more global device identifiers to each entity or device (e.g., packet engines, virtual machines, applications, services and/or other entities of each core). The aggregator may generate, provide and/or assign a unique global device identifier to each packet engine, virtual machine, application, service and/or other entity of a core. For example, the aggregator may assign, for each core of the multi-core system, a global device number to a virtual machine executing on the core. The aggregator may assign, for each core of the multi-core system, a global device number to each virtual machine of the core. Each virtual machine may include one or more packet engines processing network traffic between the one or more clients and servers. The aggregator may generate, provide and/or assign a global device identifier that is unique within and/or between cores of the multi-core system. The aggregator may generate, provide and/or assign a global device identifier that uniquely identifies an entity or device from others within a core and/or between cores of the multi-core system.

The aggregator and/or each packet engine may create, maintain and/or manage a mapping of device names (e.g., local device names) or identifiers to global device identifiers in the shared memory. Each core or packet engine may assign a local device name or identifier to each device associated with the core or packet engine, for example, virtual servers executing on the core. The local device name of an entity or device may be unique within a core. The local device name may not be unique between cores. The aggregator may use the mapping to determine a global device identifier corresponding to a local device name referenced in or identified with a data stream received from a packet engine. The aggregator may, for example, parse or extract such a local device name from the data stream.

In some embodiments, a packet engine or other device writes or transfers a data or performance stream into the shared memory. The aggregator, via an assigned thread or otherwise, may receive, poll or transfer a data stream from the shared memory, according to a schedule or predetermined interval for example. The packet engine may write a data stream from one or more devices associated with a packet engine or core. The core or packet engine may provide, assign or reference a device name (e.g., local device identifier) associated with each data stream. The core or packet engine may dump or transfer a data stream, or part thereof, to the shared memory according to a schedule or predetermined interval. In certain embodiments, the core or packet engine may provide a global device identifier associated with a data stream to the aggregator by using the device-name to global-device-identifier mapping.

In some embodiments, the aggregator determines or identifies a global device identifier for the device or entity generating the data stream. The aggregator may determine a global device identifier based on a local device identifier referenced in the data stream and/or the identity of the device or entity generating the data stream. The aggregator may associate the global device identifier with the data stream. In some embodiments, the aggregator stores the mapping and/or global device identifiers in the shared memory. The mapping may comprise any type or form of table, database or data structure for storing and/or organizing the global device identifiers. In further details of step 683, an aggregator of the multi-core system may execute a computing thread for each core of the multi-core system. An aggregator of the multi-core system may execute a computing thread on a processor or processing resource of the aggregator. In some embodiments, the aggrega-

tor executes a computing thread on a processor or processing resource assigned to the aggregator. In certain embodiments, the aggregator executes a corresponding thread in memory shared or associated with a core. The aggregator or multi-core system may provide a computing thread for each core of the multi-core system to gather data from the corresponding core. The computing thread may access shared memory to gather data or some portion of a data stream provided from a core.

The aggregator may provide a thread to operate with each active core or packet engine. The aggregator may create or execute a thread to operate with a core responsive to the core powering up, entering an active mode, or processing one or more packets. In some embodiments, the aggregator may provide a thread to operate with each entity or device of a core. A processor of the aggregator, or a processing resource assigned to the aggregator, may provide multi-threading support to handle a plurality of cores of the system. The aggregator may maintain one or more threads to handle data generated by one or more cores and/or client requests (e.g., for performance statistics or trace data).

In certain embodiments, a block of shared memory is assigned to a corresponding thread and a packet engine or core. The aggregator may allocate or assign memory to a core that is accessible by both the core and a corresponding thread. A thread may perform polling on the shared memory to detect and receive data streams (e.g., statistics of network traffic) dumped by a respective packet engine, entity or device. In some embodiments, a packet engine, entity or device informs a respective thread that a data stream is available. The thread may identify a global device identifier for a data stream, either determined by a respective packet engine, or by translating a device number associated with the packet engine or data stream using the mapping information. In further details of step 685, a first computing thread of the aggregator may collect statistics of network traffic processed by one or more packet engines on a first core. A first computing thread may generate aggregated statistics from a corresponding core by parsing the gathered data from the corresponding core. The thread may parse one or more data streams from a core or packet engine for consolidation into aggregated statistics. In some embodiments, the first computing thread may buffer, interleave and/or consolidate data streams or statistics from two or more packet engines or entities of a core. A thread may access one or more data stream dumps from the shared memory. The thread may access one or more data streams according to a schedule or predetermined interval. The thread may parse a data stream to identify one or more pieces of information. The thread may parse, sample or filter the data stream into a processed stream. The thread may extract information from a data stream to compute or determine statistics (e.g., performance and/or trace statistics) for a device of the core. In some embodiments, a thread may be assigned to handle client requests, such as a request (e.g., stat, snmpd, show, nsld) from a filtered request-response client. This thread may operate independently of any core or packet engine of the multi-core system. The thread may communicate with a client or agent, for example via an agent interface 615. In some embodiments, the thread communicates with a client or agent via a port (e.g., port 5555). The thread may access one or more data streams responsive to a client request. For example and in some embodiments, the thread may access a data stream (e.g., via shared memory) based on a global device number identified by a client or agent. The thread may parse, filter, extract or otherwise process a data stream to generate statistics responsive to the client request. This thread may execute and operate independent of other

aggregator threads and/or packet engines. In some embodiments, the thread may process data streams according to an adaptively-generated or adjusted schedule, which may depend on the processing speed of the thread and/or the transfer rate of a packet engine.

In some embodiments, a thread may be assigned to handle client requests, such as a request (e.g., dashboard request) from a streaming client. The thread and/or the aggregator may write at least a portion of a data stream into an internal buffer of the aggregator based on or responsive to the request. In some embodiments, the aggregator may generate aggregated statistics for one or more cores responsive to receiving a client request, e.g., identifying the one or more cores. For example, the aggregator may consolidate at least a portion of the statistics collected from two or more cores of the multi-core system into a buffer. In some embodiments, a data stream may be identified by a global device identifier in a client request. In certain embodiments, the thread and/or the aggregator may parse, filter, extract and/or process one or more identified data streams to generate statistics, e.g., aggregated statistics, for storing into the internal buffer.

A client or agent of the client may access the internal buffer for the requested information. In some embodiments, the thread buffers the requested information in the internal buffer for transmission (e.g., via a dashboard connection) to a dashboard, agent or client. In some embodiments, a client or agent may communicate to the thread that further information is no longer required or requested, or that the dashboard connection is closed. Responsive to this communication, the thread may halt aggregation and/or writing of data into the internal buffer.

In further details of step 687, the first computing thread may transfer the statistics with a marker to a statistics log of the multi-core system. The marker may correspond to a global device number of the first core. The first computing thread may transfer the generated statistics to a statistics log according to a schedule. The first computing thread may write collected statistics to the statistics log according to a schedule generated and/or maintained by the aggregator. The thread assigned to a core or packet engine may parse and/or process one or more data streams from the core or packet engine into aggregated statistics or aggregated data streams. In some embodiments, the thread writes or stores the processed data into a logfile or storage device (e.g., residing on the appliance or another device). The thread may, in some embodiments, write or store the processed data in an internal or local buffer of the aggregator.

In certain embodiments, the thread writes or stores each data stream unchanged or substantially unchanged into the logfile or storage device. For example, the first computing thread may transfer the collected statistics unchanged to the statistics log. The thread may transfer a data stream directly from shared memory to the logfile or storage device. Each thread may store or interleave a respective data stream into the logfile or storage device accordingly to a schedule or predetermined interval. In certain embodiments, the aggregator may create an aggregated data stream from one or more data streams, for storing into the logfile or storage device.

The multi-core system may, for example via the aggregator, maintain the statistics log for two or more cores of the multi-core system. The multi-core system may separate, partition, segment the statistics log using the markers and/or global device identifiers. The multi-core system may interleave, identify and/or mark portions of the statistics log to record or store data from different cores or devices of the multi-core system. In some embodiments, a marker or global device identifier may mark the start of data specific to a

corresponding core or device in the statistics log. In certain embodiments, a marker or global device identifier may mark the end of data specific to a corresponding core or device in the statistics log.

In some embodiments, the multi-core system provides one or more clients or their agents with access to the statistics log. The multi-core system may provide an interface, such as an applications processing interface, through which a client or agent can read or access the statistics log. In certain embodiments, the multi-core system provides the statistics log in text form. The multi-core system may provide the statistics log in any standardized format that a client can read, interpret and/or process. The aggregator may, in some embodiments, generate and/or send a notification to a first client of the one or more clients responsive to a transfer to the statistic log.

In further details of step 689, the multi-core system adaptively reschedules the transfer by monitoring the operation of each computing thread. The multi-core system may adaptively reschedule the transfer by monitoring the operation of each computing thread, core, packet engine and/or other entity of a core. For example, the aggregator may adaptively reschedule the transfer by monitoring the operation of the first core. The aggregator or the system may adaptively modify or adjust a schedule or interval for transferring data by one or more packet engines and threads. The aggregator may determine that one or more threads are operating slower than the respective packet engines.

The aggregator may determine that the dumping schedule or interval should be adjusted to synchronize operation between one or more packet engines and threads. The aggregator may determine that the logfile size is increasing at a rate beyond a predetermined threshold. The aggregator may determine or detect a change in the creation rate or transfer rate of one or more data streams. The aggregator may determine a change in the frequency of client requests. Responsive to the determination, the aggregator may adaptively modify or adjust a schedule or interval for transferring data between one or more cores, packet engines, shared memory, threads, internal buffers and/or logfiles.

A client or client agent may parse or extract particular portions of the statistics log using the markers or global device identifiers. For example, a client agent may extract statistics corresponding to a particular core or global device identifier for processing. In some embodiments, a client or client agent process the statistics log to generate aggregated statistics corresponding to one or more cores or global device identifiers. In some embodiments, responsive to a request from a client, an agent obtains data stored in the logfile for consolidation or aggregation at the client. In some embodiments, a client or agent may prefer or need consolidated or aggregated statistics for the multi-core system. The client or agent may access the logfile or storage device for one or more stored data streams. The client may process the one or more data streams, at the client, into aggregated data streams and/or statistics. In some embodiments, the client or agent may access the logfile or storage device for an aggregated data stream. The client may process the aggregated data stream, at the client, into consolidated statistics.

In some embodiments, multiple client requests and/or multiple type of client requests may occur concurrently or sequentially. A main thread of the aggregator for example, may process a request from a filtered request-response client, while another client or agent may access the logfile in connection with a nsconnmsg/nscollect request. In certain embodiments, a streaming client request (e.g., dashboard request) may be processed by providing the requested data via the internal buffer of the aggregator. The latter may be

handled by another thread of the aggregator, in parallel with the other client requests. Various threads may operate independently or in a coordinated fashion (e.g., adaptively coordinated by the aggregator in sequence).

It should be understood that the systems described above may provide multiple ones of any or each of those components and these components may be provided on either a standalone machine or, in some embodiments, on multiple machines in a distributed system. In addition, the systems and methods described above may be provided as one or more computer-readable programs or executable instructions embodied on or in one or more articles of manufacture. The article of manufacture may be a floppy disk, a hard disk, a CD-ROM, a flash memory card, a PROM, a RAM, a ROM, or a magnetic tape. In general, the computer-readable programs may be implemented in any programming language, such as LISP, PERL, C, C++, C#, PROLOG, or in any byte code language such as JAVA. The software programs or executable instructions may be stored on or in one or more articles of manufacture as object code.

While the invention has been particularly shown and described with reference to specific embodiments, it should be understood by those skilled in the art that various changes in form and detail may be made therein without departing from the spirit and scope of the invention as disclosed herein.

We claim:

1. A method comprising:

- (a) maintaining, in shared memory of a multi-core system intermediary between one or more clients and servers, a global device number for each core of the multi-core system, each core comprising one or more packet engines processing network traffic between the one or more clients and servers;
- (b) executing, by an aggregator of the multi-core system, a computing thread for each core of the multi-core system;
- (c) collecting, by a first computing thread of the aggregator, first statistics of network traffic processed by one or more packet engines on a first core and collecting, by a second computing thread of the aggregator, second statistics of network traffic processed by one or more packet engines on a second core; and
- (d) maintaining, by the multi-core system, a statistics log for two or more cores of the multi-core system, including at least the first core and the second core, transferring, by the first computing thread, the first statistics with a marker to the statistics log, the marker corresponding to a global device number of the first core.

2. The method of claim 1, wherein (a) further comprises assigning, for each core of the multi-core system, a global device number to a virtual machine executing on the core.

3. The method of claim 1, wherein (a) further comprises assigning, for each core of the multi-core system, a global device number to each virtual machine of the core, each virtual machine comprising one or more packet engines processing network traffic between the one or more clients and servers.

4. The method of claim 1, wherein (c) further comprises consolidating, by the aggregator, at least a portion of statistics collected from two or more cores of the multi-core system into a buffer.

5. The method of claim 1, wherein (d) further comprises writing the collected statistics to the statistics log according to a schedule of the aggregator.

6. The method of claim 5, wherein (d) further comprises adaptively rescheduling, by the aggregator, the transfer by monitoring the operation of the first core.

7. The method of claim 1, wherein (d) further comprises transferring the collected statistics unchanged to the statistics log.

8. The method of claim 1, wherein (d) further comprises transferring, by the second computing thread, the second statistics with a second marker to the statistics log, the second marker corresponding to a global device number of the second core.

9. The method of claim 1, further comprising sending, by the aggregator, a notification to a first client of the one or more clients responsive to the transfer.

10. The method of claim 1, further comprising providing, by the multi-core system to the one or more clients, access to the statistics log.

11. A system comprising:

an intermediary device intermediary between one or more clients and servers;

shared memory between multiple cores of the intermediary device, the shared memory maintaining a global device number for each of the multiple cores, each core comprising one or more packet engines processing network traffic between the one or more clients and servers; and an aggregator of the intermediary device, executing a computing thread for each of the multiple cores, comprising at least a first computing thread collecting first statistics of network traffic processed by one or more packet engines on a first core and a second computing thread collecting second statistics of network traffic processed by one or more packet engines on a second core, and transferring the first statistics with a marker to a statistics log of the multi-core intermediary device, the marker corresponding to a global device number of the first core, the statistics log maintained for two or more cores of the multi-core intermediary device, including at least the first core and the second core.

12. The system of claim 11, wherein the aggregator assigns a global device number to a virtual machine executing on each core.

13. The system of claim 12, wherein the aggregator assigns a global device number to each virtual machine of each core, each virtual machine comprising one or more packet engines processing network traffic between the one or more clients and servers.

14. The system of claim 11, wherein the aggregator further consolidates at least a portion of statistics collected from two or more cores of the multi-core intermediary device into a buffer.

15. The system of claim 11, wherein the first computing thread writes the collected statistics to the statistics log according to a schedule of the aggregator.

16. The system of claim 15, wherein the aggregator adaptively reschedules the transfer by monitoring the operation of the first core.

17. The system of claim 11, wherein the first computing thread transfers the collected statistics unchanged to the statistics log.

18. The system of claim 11, wherein the aggregator transfers the second statistics with a second marker to the statistics log, the second marker corresponding to a global device number of the second core.

19. The system of claim 11, wherein the aggregator sends a notification to a first client of the one or more clients responsive to the transfer.

20. The system of claim 11, wherein the aggregator provides the one or more clients access to the statistics log.